ABSTRACT

Properties of minimum spanning trees and fractional quantum Hall states

Thomas Sundal Jackson

2010

This dissertation consists of work done on two disjoint problems. In the first two chapters I discuss fractal properties of average-case solutions to the random minimal spanning tree (MST) problem: given a graph with costs on the edges, the MST is the spanning tree minimizing the sum of the total cost of the chosen edges. In the random version the costs are quenched random variables.

I solve the random MST problem on the Bethe lattice with appropriate boundary conditions and use the results to infer fractal dimensions in the mean-field approximation. I find that connected components of the MST in a window have dimension $D = 6$, which establishes the upper critical dimension $d_c = 6$. This contradicts a value $d_c = 8$ proposed previously in the literature; I correct the argument that led to this value. I then develop an exact low-density expansion for the random MST on a finite graph and use it to develop an expansion for the MST on critical percolation clusters. I prove this perturbation expansion is renormalizable around $d_c = 6$. Using a renormalization-group approach, I calculate the fractal dimension D_{p} of paths on the latter MST to first order in $\varepsilon = 6 - d$ for $d \leq 6$, with the result $D_{\mathrm{p}} \sim 2 - \varepsilon/7$.

In the final chapter, I investigate the correspondence between wavefunctions in the fractional quantum Hall effect obtained as blocks of a conformal field theory (CFT) versus those defined as zero-energy eigenstates of projection Hamiltonians, specifically one which forbids three particles to come together in one of two linearly-independent states of relative angular momentum six and all states of lesser relative angular momentum. I construct zero-energy states from amplitudes of superconformal currents using a result due to Simon. The counting of edge excitations of these states agrees with the character formula for the superconformal Kac vacuum module at generic central charge c, which implies this Hamiltonian is gapless for all c. I attempt to obtain a rational theory by projecting out additional states, focusing on the $M(3, 8)$ and tricritical Ising CFTs.

Properties of minimum spanning trees and fractional quantum Hall states

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Thomas Sundal Jackson

Dissertation Director: Nicholas Read

May 2010

Contents

1 Introduction **1**

1.1 Geometric properties of minimum spanning trees 1

1.1.1 Context and motivation 1

1.1.2 Summary of prior work 6

1.1.3 Outline of new results 7

1.1.4 Future directions 11

1.2 Projection Hamiltonians in the fractional quantum Hall effect 11

1.2.1 Context and motivation 12

1.2.2 Summary of prior work 13

1.2.3 Outline of new results 15

1.2.4 Future directions 16

2 Minimum spanning trees: Bethe lattice **17**

2.1 Minimum spanning trees, percolation and spin glasses 18

2.1.1 Minimum spanning trees (MSTs): properties and algorithms 19

2.1.2 Bond and invasion percolation 22

2.1.3 MST algorithms and percolation 26

2.1.4 Spin glasses 27

2.1.5 Strongly-disordered spin glass model and MSTs 29

2.1.6 Ground states and critical dimension of the NS model 32

2.2 The MST problem on the wired Bethe lattice 36

2.2.1 Definitions . . . 37
2.2.2 Percolation on the Bethe lattice . . . 38
2.2.3 MST connectedness functions . . . 39
2.2.4 Asymptotic behavior of the iterated kernel . . . 45
2.2.5 Asymptotics of the two-point function and mass . . . 50
2.2.6 k-point connectedness functions and moments of the mass . . . 54
2.2.7 Fractal dimensions and mean-field theory . . . 56
2.3 Related applications . . . 58
2.3.1 Random, locally treelike graphs . . . 58
2.3.2 Poisson-weighted infinite tree . . . 59
2.4 Conclusions . . . 60

3 Minimum spanning trees: Loop corrections 62
3.1 Graphical expansions for percolation . . . 64
3.1.1 Two-point connectedness functions . . . 64
3.1.2 n-point connectedness functions . . . 68
3.1.3 Equivalence with the Potts model . . . 69
3.1.4 Topological properties of the diagrammatic weights . . . 72
3.2 A graphical expansion for MSF paths . . . 74
3.2.1 MST paths as geodesics . . . 74
3.2.2 Identifying MST paths through binary comparisons . . . 77
3.2.3 A lattice expansion for MSF paths . . . 79
3.2.4 Expansion in terms of topological graphs . . . 83
3.3 A continuum perturbation expansion for MSF paths . . . 86
3.3.1 The excluded volume constraint . . . 87
3.3.2 Continuum expansion for MSF paths . . . 88
3.3.3 Scaling and fractal dimensions: mean-field theory revisited . . . 94
3.3.4 Beyond mean-field theory: breakdown of perturbation theory for $d < 6$ . . . 95
3.4 Renormalizability of the MSF field theory . . . 101

3.4.1 Definitions . . . 103
3.4.2 Effect of the $\mathcal{O}_{\text{MSF}}$ operator . . . 105
3.4.3 Power counting for MSF Feynman integrals . . . 107
3.4.4 Factorization properties of MSF diagrammatic weights . . . 110
3.4.5 Proof of renormalizability . . . 117
3.4.6 Cancellation proof for subleading terms . . . 124
3.5 One-loop renormalization group calculation . . . 127
3.6 Postscript: remarks for $p > p_c$ . . . 134
3.6.1 The "superhighways" conjecture . . . 134
3.7 Conclusions . . . 136

4 Fractional quantum Hall wavefunctions 139
4.1 Introduction to the fractional quantum Hall effect . . . 142
4.1.1 Integer and fractional effects . . . 142
4.1.2 Lowest Landau level polynomials . . . 145
4.1.3 Laughlin states and pseudopotentials . . . 146
4.1.4 FQHE states, Chern-Simons theory and conformal blocks . . . 150
4.2 Summary of conformal field theory . . . 151
4.2.1 Conformal invariance . . . 151
4.2.2 Radial quantization . . . 155
4.2.3 Minimal models . . . 159
4.2.4 TQFT, RCFT and all that . . . 161
4.2.5 Superconformal symmetry . . . 165
4.3 Clustered states, residues and filtrations . . . 166
4.3.1 k-body pseudopotential formalism . . . 166
4.3.2 The filtration method . . . 172
4.3.3 Filtrations and symmetric polynomials . . . 174
4.3.4 Superconformal theories at generic c . . . 179
4.3.5 The $M(3, p)$ series . . . 187

4.3.6 Remarks on the tricritical Ising model 192
4.4 Discussion 199

List of Figures

1.1 An example of the minimal spanning tree (MST) problem. The figure on the left shows the input, a graph with costs associated to the edges (the numbers). On the right is the MST of the graph given those costs: the subset of edges (drawn with thick lines) which minimizes the total cost of selected edges while connecting all vertices to each other. 2

1.2 Depiction of a path between two points (the open circles) on one realization of the random MST. The fractal dimension D_{p} measures how the fractal mass of the path (number of thick edges) scales with the Euclidean separation of the endpoints (the gray line). 4

1.3 Illustration of the wired random MST construction arising in the strongly disordered spin glass model of Newman and Stein [23]. Within any "window" or finite region in the interior (gray line), the MST on an infinite graph with wired boundary conditions will appear to be a minimal spanning forest of several space-filling, disjoint connected components. One such component is highlighted in black. . . . 5

1.4 Numerically determined results for the fractal dimension of paths on the MST obtained in the literature. 6

2.1 A sample realization of the MST (solid lines) on the wired Cayley tree. The graph shown has $\sigma = 2$ and radius $M = 5$. Vertices on the boundary are connected with boundary edges of vanishing cost, shown here as thin lines. All vertices in the interior are connected to the boundary only once. 38

2.2 Construction of a path on the Bethe lattice starting from the origin (labeled m), passing through the vertex labeled 0, and connecting to the boundary, here shown as a shaded line. For this path to lie on the MST, the edge costs must obey the set of inequalities given in eq. (2.16). Note that each vertical line (both solid and dashed) stands for a subtree, not just a single edge, so this diagram depicts the entire Bethe lattice. 41

2.3 The set of paths contributing to $C^{(2)}(m; p_0)$, drawn using the same conventions as figure 2.2. Here I sum over all $m_1 \geq 0$, $m_2 \geq 0$ such that $m_1 + m_2 = m$. 44

3.1 Perturbation expansion for the path vertex function given in (3.76). I represent the path vertex by the open circle connected to a wavy line. Note that these diagrams do not include any notation corresponding to the operation of $\mathcal{O}_{\text{MSF}}$. The lowest non-trivial term is the second one on the right-hand side. 98

3.2 Depiction of an arbitrary graph contributing to the MSF path connectedness function. The various root vertices used in the definition of this function are labeled with open circles. 103

3.3 Factorization properties of $d_{\text{MSF}}(\mathcal{G}|\pi)$ for connected subgraphs that are proved in the text for orderings in the class $\pi_{[0]}$. 114

3.4 The factorization property of $d_{\text{MSF}}(\mathcal{G}|\pi)$ proved in the text for arbitrary ordering of edge costs. 115

3.5 The 1PI one-loop diagrams for vertex functions that are ultraviolet divergent in six dimensions: a) self-energy, b) cubic interaction vertex function, c) mass-insertion vertex function, d) MSF path vertex function. 130

4.1 Tabulation of the dimensions $D_{k,r}$ of the space of symmetric polynomials in k variables that are homogenous of degree r, and of the subspace of those polynomials which are translationally invariant, which has dimension $\tilde{D}_{k,r}$. 170

4.2 Example of various partitions involved in defining translationally-invariant symmetric polynomials, here in $k+1=5$ variables. Parts $a), b)$ illustrate the first case discussed in the text, in which the new partition $w(\lambda) = \lambda'$ is the transpose of the admissible partition λ. Parts $c), d)$ and $e)$ illustrate the extra rule we adopt when the width of λ is equal to the number of variables. 175

4.3 Schematic illustration of the pattern of connections involved in the various symmetric polynomials. Here the particle coordinates occupy boxes of a Young diagram and each diagonal or vertical line stands for a factor of $(z_{1,i} - z_{2,j})$. Figure $a)$ depicts $\vartheta_{\{3,1,1\}}$ for $k=5$ particles. Figures $b)$ and $c)$ show $\overline{\vartheta}_{\{3,1,1\}}$ for the case of rows of equal and unequal width, respectively. 177

4.4 Diagram of the organization of factors in the state (4.97). Part $a)$ labels the organization of the coordinates, each of which occupies a box in the diagram of κ. Part $b)$ depicts each product appearing in (4.97): here for clarity I only draw the first factor in each product, using thin lines to stand for factors of coordinate differences $(z^{(\alpha)}_{i,j} - z^{(\alpha')}_{i',j'})$. 185

4.5 Diagram of the organization of factors in the state (4.109), which should be compared with figure 4.4 showing the analogous situation for the generic superconformal theory. Here for clarity I only draw the first factor in each product, using thin lines to stand for factors of coordinate differences $(z^{(\alpha)}_{i,j} - z^{(\alpha')}_{i',j'})$. 190

Acknowledgements

Success is the ability to go from one failure to another with no loss of enthusiasm.
— Winston Churchill

I have been dependent on the support of many people while performing the work which led to this dissertation, and there are too many of them to adequately list in the space of a page.

I'd like to thank my advisor, Nicholas Read, for his guidance and continued patience during our collaboration; he has set a standard to which I hope to hold myself, my research and my teaching. I'd also like to thank Steve Girvin, Simon Mochrie and Walter Goldberger, for investing the time to serve on my thesis committee and for helpful comments and criticism on a previous draft. I am grateful for the financial support provided by the National Science Foundation and the Warren A. and Hibernia S. Tyrell fellowship in physics.

I'd like to thank my colleagues both in condensed matter theory — Doron Bergman, Lev Bishop, Benoit Estienne, Ludovic Jaubert, Chris Laumann, George Mias, Sid Parameswaran, Stephen Powell and Terri Yu, and outside of it — Caitie Barrett, Christopher Crick, Erik Encarnacion, Adrienne Erickcek, Catherine Flynn, Daniel Grin, Rajini Haraksingh, Kendall Heitzman, Zigmund Kermish, Lauren Lippiello, Jeff Mankoff, Molly A. Martinez, Claire McMurray, Sasha Santee, Karl Waern and Charlotte Walker — for their advice and camaraderie over the years.

And last but certainly not least, I thank my parents for supporting me prior to graduate school and for valuing education so highly.

Chapter 1

Introduction

The work presented in this dissertation takes the form of two disjoint projects undertaken with my advisor, Nicholas Read: one on the average-case geometric properties of minimum spanning trees (MSTs) [1, 2] and another (still in progress) on properties of trial ground states in the fractional quantum Hall effect (FQHE) obtained as zero-energy eigenstates of special projection Hamiltonians. In this chapter, I discuss each project in turn, attempt to place it in context, present its motivation and summarize the new results given in this dissertation.

1.1 Geometric properties of minimum spanning trees

The two chapters devoted to the minimal spanning tree problem are based on the papers [1] and [2], and follow their organization closely. These sources were written collaboratively by Nicholas Read and myself.

1.1.1 Context and motivation

Minimal spanning trees in graph theory, probability and computer science

The minimum spanning tree problem (MST) is an elementary, well-studied problem in computer science in the particular field of combinatorial optimization [3–7]. The problem is defined on an undirected, connected graph with costs ℓ_e associated with each edge e. I define a spanning

tree as a subset of the graph's edges which connects all the vertices to each other (*i.e.*, spans them) without containing any cycles of edges (*i.e.*, it is a tree in the graph theory sense). At least one such tree exists if we assume the graph is connected. Note that the absence of cycles implies that there is a unique path between any two points on the tree.

A *minimum* spanning tree is a spanning tree which minimizes the total cost of edges selected to be on the tree, and the MST problem is to construct this tree given the graph and costs as input. I give an example of a MST in figure 1.1. For a given input, this is rather simple to do computationally, and can be accomplished by any one of a number of straightforward algorithms [5, 8, 9].

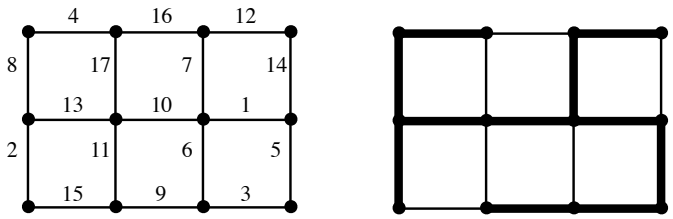

Figure 1.1: An example of the minimal spanning tree (MST) problem. The figure on the left shows the input, a graph with costs associated to the edges (the numbers). On the right is the MST of the graph given those costs: the subset of edges (drawn with thick lines) which minimizes the total cost of selected edges while connecting all vertices to each other.

Minimum spanning trees arise in many situations of network design. This is because they solve what I refer to as the strong-disorder limit of the optimal path problem: suppose you have some commodity that you need to transport between two points via some network (the graph), with a particular cost for using each segment (edge) of the network. Here the quantity to be minimized is the sum of the costs of edges in a path with specified endpoints. What I refer to as the "strong disorder limit" below occurs when this sum may be approximated by its largest term — *i.e.*, when transport is dominated by bottlenecks. Then the minimal spanning tree is the union of the optimal paths between *any* pair of points on the graph [10–13].

One specific and ubiquitous application of minimal spanning trees is the IEEE 802.1D standard [14], which governs the behavior of network bridges connecting segments of an Ethernet computer network. This standard calls for the bridges to collectively compute (an object related to) the minimal spanning tree of the graph describing their connectivity, resulting in a unique path for

data to be passed between any pair of bridges, which greatly simplifies the routing of data in the network.

Finally, a frequently encountered specialization of the MST problem is the continuum MST problem [7, 15]. Here the vertices of the graph are points embedded in some metric space, the graph is the complete graph (with an edge connecting any pair of vertices), and the edge cost is the distance between its endpoints given by this metric. This has applications in the study of clustering (or classification) of data: given a set of data points we would like to extract sets of these points which are "similar" or "close together" in a sense captured by the metric. One computationally inexpensive way to accomplish this is to compute the continuum MST of these points, and then remove the longest edges according to some criterion, resulting in a partition of the vertices into a set of connected components (the clusters) [16].

Minimal spanning trees in physics

Minimal spanning trees have also found applications in more physical situations. Here I find it useful to define a variant of the MST problem which is amenable to treatment by statistical mechanics methods: this is the specific problem I study in chapters 2 and 3. I fix the underlying graph G to be a portion of some regular lattice (in section 2.2 I formulate mean-field theory on the Bethe lattice and then extend the results to Euclidean lattices of high dimension in section 2.2.7). In this dissertation I concern myself with long-distance properties of the MST, so by the standard universality arguments I expect my results to depend only on the lattice dimensionality d, independent of the microscopic lattice structure.

I then take the edge costs to be independent, identically distributed (iid) random variables from some continuous probability distribution (as we will see in section 2.1.1, the shape of this distribution will not be important). Because the edge costs are drawn from a continuous distribution, they are distinct with probability one, and the MST produced from a fixed realization of the edge costs is unique. I then perform a *quenched* average over the edge costs: *i.e.*, I obtain a probability distribution on the set of spanning trees by generating many realizations of edge costs, solving the MST problem on each one, and constructing a histogram of the number of times each spanning tree

is the MST. (This is how one would set up a numerical calculation of the properties given below, but it is important to emphasize that my results are entirely analytic and obtained by other means.)

This random MST problem (which I refer to simply as "the MST problem" in chapters 2 and 3) has applications to the problem of transport in highly disordered media, again via the connection with the strongly disordered optimal paths. These were studied in [17–19]. Two main applications that have been examined in the literature are undirected polymers in disordered media [11, 19] and patterns of current flow in random resistor networks [20–22].

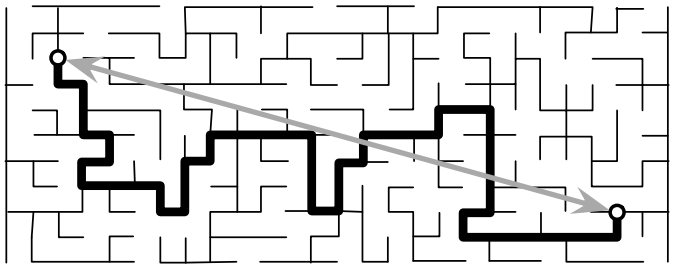

Figure 1.2: Depiction of a path between two points (the open circles) on one realization of the random MST. The fractal dimension D_{p} measures how the fractal mass of the path (number of thick edges) scales with the Euclidean separation of the endpoints (the gray line).

A noteworthy use of minimal spanning trees which does not arise from the study of optimal paths is the spin glass model proposed by Newman and Stein [23]. In the strong disorder limit, the problem of finding the ground state of an Ising spin glass may be directly mapped on to the MST problem (see section 2.1.5) — specifically, the random MST problem with edge costs on the boundary of the graph taken to be much smaller than any edge cost in the interior. This means that any solution of the random MST problem on this restricted ensemble will consist of a minimal spanning *forest* (MSF): the MST will consist of a number N_{MSF} of disjoint trees in the interior connected by the boundary edges, as shown in figure 1.3. In the Newman and Stein model, each component of the MSF corresponds to a "frozen" cluster of spins whose relative orientations are fixed in all ground states, and whose absolute orientation is fixed by the boundary conditions. The number of MSF components is therefore related to the number of ground states of this model, which is a topic of controversy in the study of short-ranged spin glass models (see section 2.1.4). The scaling of this number also sets the upper critical dimension of the random MST model (see

section 2.1.2), and hence the domain of validity of mean-field theory.

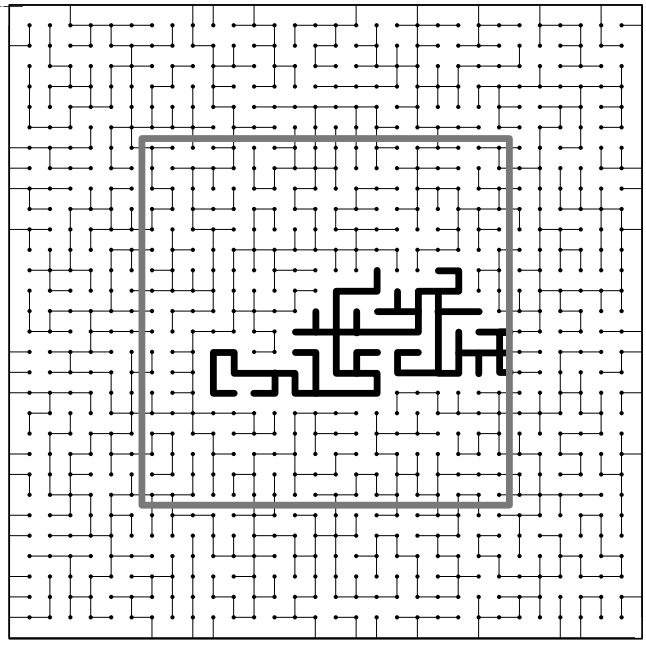

Figure 1.3: Illustration of the wired random MST construction arising in the strongly disordered spin glass model of Newman and Stein [23]. Within any "window" or finite region in the interior (gray line), the MST on an infinite graph with wired boundary conditions will appear to be a minimal spanning forest of several space-filling, disjoint connected components. One such component is highlighted in black.

From the Newman-Stein model, we may come full circle and return to the domain of computer science, since questions about the ground state structure of spin glasses have direct application to the existence, multiplicity and organization of solutions to hard optimization problems in computer science [24–26].

The applications above provide the motivation for studying the quantities which will be of central interest in chapters 2 and 3: namely the fractal dimension D_{p} of paths on the MST, and the number N_{MSF} of connected components of the MSF induced in a window by a MST with wired boundary conditions, or equivalently the fractal dimension D_{MSF} of one such connected component (since the MSF components are space-filling).

1.1.2 Summary of prior work

The minimal spanning tree problem was first considered by Borukva [27] in 1926, who also gave the first algorithm for its solution. It is beyond the scope of this dissertation to summarize the work done on this problem in the computer science literature; I refer the reader to [28] for a historical overview and to [29] for a readable summary of the major algorithms.

The strong disorder optimal path problem was originally considered by Cieplak, Maritan and Banavar [11]. The connection between the strong disorder optimal path problem and MSTs was first noticed in [17] and studied further in [10, 12, 20, 22, 30, 31] (note that not all of these deal with Euclidean lattices). As mentioned above, applications to polymer physics were discussed in [11, 19] and to random resistor networks were given in [20–22, 32].

There exist in the literature several numerical simulations [10, 11, 18, 19, 31, 33] determining D_{p}, the fractal dimension of the strong disorder optimal path. These results are summarized in table 1.4. To my knowledge no numerical studies of D_{MSF} exist.

d	D_{p}
1	1
2	1.22 ± 0.01 [10, 11, 31]; 1.218 ± 0.001 [33]; 1.222 ± 0.003 [18]
3	1.42 ± 0.02 [10, 11]; 1.44 ± 0.02 [31]; 1.43 ± 0.03 [19]
4	1.59 ± 0.02 [11]

Figure 1.4: Numerically determined results for the fractal dimension of paths on the MST obtained in the literature.

Newman and Stein's strongly disordered spin glass model was proposed in [23]. Implications of this model for spin glasses with finite disorder strength were considered in [34]. Some discussions of similar issues in finite-range spin glasses are given in, *e.g.*, [35–37]. Related considerations also arose in the quantum spin glass (or random transverse field Ising model) [38].

To my knowledge, there has been little *analytical* work done on the random MST problem. An (incorrect; see section 2.1.6) conjecture for the upper critical dimension was made in [23], and in [21] the scaling properties of the costs were studied and related to scaling exponents in percolation. As far as I am aware, this is also the only prior use of the connection between Kruskal's algorithm for computing the MST [8] and bond percolation [39, 40] which forms the basis for all

my calculations (see section 2.1.3).

Spanning tree models of a different sort than those I consider here have been studied in the probability community. The question of the number of connected components was raised in a construction of a minimum spanning forest in the continuum model directly in infinite volume by Aldous and Steele [7, 15], who suggested that the forest has a single connected component in all dimensions d. It has been proved that there is a single connected component for this model in $d = 2$ [41], but for larger d there has not so far been agreement even at a heuristic level.

Although beyond the scope of this thesis, I also mention in passing the recent work on the scaling limit of MSTs in two dimensions [33, 42, 43]. Interest in this problem stems in part from the numerical observation in [33, 42] that the MST problem in two dimensions is one of the few known examples of a system which is scale invariant but not conformally invariant.

1.1.3 Outline of new results

I now summarize the original results obtained in chapters 2 and 3, which recapitulate the material of [1] and [2] respectively. In these chapters I develop an analytical theory for the random MST problem which takes a form similar to those for critical phenomena in, for example, Ising spin systems, or percolation. This means that the results presented in this dissertation fall into two main parts: in chapter 2 I first define a mean-field theory, whose predictions are valid for d sufficiently large. In section 3.3.4, a perturbation analysis around the mean-field theory will reveal that it breaks down below an upper critical dimension d_c (inferred by geometric means in section 2.2.7), and the perturbation theory can be traded in for an asymptotic expansion in powers of $\varepsilon = d_c - d$ for $d < d_c$, calculated by standard methods of statistical field theory. The derivation of this expansion and the proof (to all orders in perturbation theory) of its renormalizability is the central topic of chapter 3.

Bethe lattice and mean-field theory

In chapter 2 I consider the MST problem on the Bethe lattice, which motivates the development of my mean-field theory. This model can be solved exactly in some cases, or exactly for asymptotic,

universal properties in more cases. In section 2.3.2, I also investigate the related Poisson-weighted infinite tree (PWIT) model [15, 44, 45], which can be obtained as a limit of the Bethe lattice as the coordination number goes to infinity and defines the mean-field theory of the continuum model defined above. As motivated in the previous section, the main objects I will be concerned with are D_{p}, the fractal dimension of paths on the MST, and D_{MSF}, the fractal dimension of connected components of the MST visible within a window. This latter quantity is important because it sets the upper critical dimension d_c for my model, above which mean-field theory is exact.

The Bethe lattice is, of course, itself a tree, so attention to the definition of boundary conditions in the spirit of [23] is essential in producing nontrivial behavior. Specifically, in section 2.2.1 I adopt wired boundary conditions which mean that the minimal object is a spanning *forest* (MSF): a spanning, vertex-disjoint collection of trees. The MSF will have a nontrivial number N_{MSF} of connected components, whose behavior is the central issue for the Bethe lattice model, for reasons given in section 2.1.6. As the size of the lattice goes to infinity, the statistical properties of the MSF have a well-defined limit, and I then calculate the number $M(m)$ of vertices that are connected to the central site and lie within m steps, which I call the "mass" within m steps. I find (in section 2.2.5) that it scales as

$$M(m) \sim m^3 \tag{1.1}$$

as $m \to \infty$. The standard method for transferring results from the Bethe lattice to a Euclidean lattice (see section 2.2.7) entails that distance m on the Bethe lattice corresponds to distance squared, $m \sim R^2$ on the Euclidean lattice. Using this relation, I find that the mass within distance R of the origin scales as

$$M(R) \sim R^6, \tag{1.2}$$

so that the tree has fractal dimension $D_{\text{MSF}} = 6$. As the trees fill the lattice, this means that the number of connected components that intersect a ball of radius R, denoted $N(R)$, scales as

$$N(R) \sim R^{\#_{\text{MSF}}} \sim R^{d-6}, \tag{1.3}$$

as $R \to \infty$, so the tree "proliferation exponent" [39, 46] $\#_{\text{MSF}} = d - D_{\text{MSF}} = d - 6$. Note that on

the Bethe lattice, $N(m)$ increases exponentially with m, but there is a power-law correction factor of m^{-3} which produces the behavior relevant for the Euclidean lattice. By similar arguments I find that the mass of an MST path within distance R of the origin scales as

$$M_p(R) \sim R^{D_{\mathrm{p}}} = R^2, \tag{1.4}$$

so that, in the mean-field approximation, paths on the MST behave as random walks. Results (1.3) and (1.4) constitute the main qualitative results of chapter 2.

The proliferation exponent $\#_{\mathrm{MSF}}$ cannot become negative, so even if the mean field theory is indeed valid in high dimensions, it must break down in sufficiently low dimensions. The non-zero proliferation exponent $\#$ is the geometric counterpart to the violation of hyperscaling relations in critical phenomena at $d > d_c$ [39, 46], that is hyperscaling is obeyed when $\# = 0$. Because the union of the MSF clusters is space-filling, I must have

$$D_{\mathrm{MSF}} = \begin{cases} 6 & (d \geq 6), \\ d & (d \leq 6); \end{cases} \tag{1.5}$$

$$\#_{\mathrm{MSF}} = \begin{cases} d-6 & (d \geq 6), \\ 0 & (d \leq 6). \end{cases} \tag{1.6}$$

This result contrasts with that of Newman and Stein, who suggested that $d_c = 8$ for MSTs [23], in the sense of nonzero $\#_{\mathrm{MSF}}$ defined above. I correct their argument in section 2.1.6. Finally, I note that these values differ from known results for uniform spanning trees (USTs) where each spanning tree on a finite graph is given equal probability, namely $d_{c\,\mathrm{UST}} = 4$ and $D_{\mathrm{UST}} = 4$ for $d > 4$ [47], so this establishes that the random MST model is in a distinct universality class from USTs (at least for $d > 4$) [42, 48].

Perturbative corrections to mean-field theory below d_c

In chapter 3 I formulate a continuum perturbation expansion for the geometry of the MST in Euclidean space, which is close to being a field theory for MSTs. As with the mean-field theory,

it is based on the connection between bond percolation and Kruskal's algorithm (section 2.1.3). This formulation makes it possible first to rewrite the essentials of the Bethe lattice solution in a form similar to a mean-field theory, and a corresponding calculation of the main correlation function. The terms omitted from the mean field theory in Euclidean space can then be studied in a perturbation expansion, which has no long-distance (infrared) divergences when $d > d_c = 6$. This justifies *post hoc* the validity of the mean field theory for exponents in $d \geq 6$, at least within perturbation theory. Because $D_{\text{MST}} = d$ for $d < d_c$, in this chapter I largely restrict my attention to the fractal dimension of paths D_{p}, which exhibits non-trivial behavior. For $d \leq 6$, it becomes technically difficult to construct the MST on the entire lattice, so as a simplification I consider the MST problem on critical percolation clusters. The implications of this restriction are discussed in section 3.6.

In section 3.2 I obtain a lattice expansion for connectedness functions on components of the MSF. Although the primary utility of this result is to provide a basis for the continuum field theory of section 3.3, the lattice expansion is of interest in its own right and could be taken as the starting point for a mathematically rigorous formulation of these results using *e.g.*, lace expansion methods [49].

In section 3.3 I obtain a perturbative expansion in terms of Feynman diagrams from the continuum limit of this lattice expansion. I prove in section 3.4 that this expansion may be renormalized in a consistent manner to all orders. I emphasize that both these results are technically non-trivial, due to the non-locality of the MST problem. My diagrammatic expansion is defined in terms of performing various manipulations on percolation diagrams, rather than being generated by an action functional, and it appears likely that no such action defined in terms of a finite number of locally-interacting fields exists (I will elaborate on this point in section 3.7). The fact that this non-local optimization problem is, nonetheless, sufficiently local to permit the use of renormalization group methods is perhaps the most noteworthy result of the chapter.

In section 3.5 I then perform a one-loop renormalization group calculation and obtain the fractal

dimension of paths on the percolation cluster MST as an asymptotic series in $\varepsilon = d_c - d$, namely

$$D_{\mathrm{p}} \sim 2 - \varepsilon/7 + \mathcal{O}(\varepsilon^2) \tag{1.7}$$

as $\varepsilon \to 0$. I argue that this result also gives the dominant scaling behavior for widely separated points on percolation clusters above the percolation threshold, for MSTs and the continuum model; the results of [19] provide empirical support for this argument. There do not appear to be any scaling relations that relate the geometric exponents for MSTs to those for percolation, unlike those found for the costs in [21], even though the critical dimension $d_c = 6$ is the same.

Equation (1.7) is the main quantitative result of chapter 3. Because of the dangers involved in setting, say, $\varepsilon = 2$ in an asymptotic expansion and the influence of the subdominant effects I have neglected for finite sample sizes, I do not feel a meaningful comparison of the results presented in this dissertation with the numerical results of [10, 11, 33, 42] and table 1.4 is possible.

1.1.4 Future directions

My calculations can in principle be extended to other exponents, such as those defined in [42], or carried to higher orders in ε. They can also be extended to include statistical properties that involve the cost of the MST. Finally, I may finish the extension of the results of chapter 3 to supercritical percolation clusters and address the "superhighways" conjecture discussed in section 3.6.1. A more relevant question is whether the methods I have developed are applicable to other, potentially more interesting problems, such as different models of short-ranged spin glasses.

1.2 Projection Hamiltonians in the fractional quantum Hall effect

The final chapter describes the results of a completely different research project, done in collaboration with Nick Read and Steve Simon. This work is still in progress and is unpublished at the time of writing.

1.2.1 Context and motivation

The discovery and explanation of the integer and fractional quantum Hall effects is one of the major achievements of modern condensed matter physics, representing the first experimental realization of topologically ordered states of matter. The description of these topological phases remains an open problem and the subject of much current activity (see next section).

In this part of the dissertation, I will discuss certain issues which arise in describing states of the fractional quantum Hall effect (FQHE). Since the work of Laughlin [50], the main tool we have for describing phases of the FQHE has been the formalism of trial wavefunctions, in which a microscopic, first-quantized electron wavefunction is given which is taken to be representative of the entire phase: the physical electron wavefunction is assumed to be adiabatically connected to the trial wavefunctions, whose validity may be established by numerical exact diagonalization studies done on systems of several particles.

In the literature there are two main ways to describe a trial wavefunction: either as a block of some conformal field theory (CFT), or as the zero-energy eigenstate of a specially-chosen projection Hamiltonian. Both approaches have strengths and weaknesses: with the former definition, topologically-invariant defining characteristics of the phase such as braiding statistics of quasihole excitations (see section 4.1.4) may easily be read off from the field content of the CFT, while states defined using the latter definition are easier to realize explicitly and work with numerically. The overarching problem that my work in this chapter is directed at is that of reconciling the two approaches to describing wavefunctions.

The results given in this chapter concern a very specific instance of this broad correspondence we would like to establish: I consider a projection Hamiltonian giving rise to a clustered Hall state [51], which is so called because the Hamiltonian forces the trial wavefunction to vanish as some power of the particle separation whenever some number of particles are brought to the same spot. Interactions of this type, while artificial, can be thought of as some sort of dressed Coulomb repulsion projected to the lowest Landau level. The specific Hamiltonian studied in this chapter forbids one type of behavior as three particles come together with relative angular momentum six (in units where $\hbar$ has been set equal to one); this is significant because there is a two-dimensional space of

ways this can happen; the allowed behavior is selected by free parameter in the Hamiltonian.

This situation offers an opportunity to formulate a connection between projection Hamiltonians and CFTs: Simon [52] recently evaluated the amplitude of an arbitrary number of supercurrents in an arbitrary conformal field theory with supersymmetry (a superconformal field theory, or SCFT), and found that it took on a similar clustered form. This suggested that the FQHE states produced by the Hamiltonian described in the previous paragraph might be generated by superconformal field theories, with the free parameter in the Hamiltonian setting the central charge c of the theory.

Simon found that his result correctly describes the angular-momentum ground state for any c. In section 4.3.4 we construct all excited states allowed by the Hamiltonian and find that the resulting basis is, in fact, independent of c: although the basis functions themselves change, there are still the same number of linearly-independent functions at any level of angular momentum. This is confirmed by counting the states: the character thus obtained is the character of the CFT describing chiral edge excitations of the state, and for this basis the character is found to agree with that of a non-minimal SCFT at generic central charge.

Stable quantum Hall phases are described by unitary CFTs [53], and obtaining a unitary minimal theory requires excluding additional states in accordance with the singular vector structure of the SCFT. I am able to construct basis states and find the additional projection operator needed to obtain the (non-unitary) $SM(2,8)$ SCFT. The simplest unitary minimal SCFT appears to be the tricritical Ising model, $SM(3,5)$, but here I have been able to make less progress. The necessary projection operator here appears to require seven-particle interactions.

1.2.2 Summary of prior work

The literature on the definition, study and classification of trial wavefunctions in the FQHE is vast; this review is necessarily quite narrow and focuses on those methods which attempt to study the polynomial properties of the trial wavefunctions directly.

Wen and collaborators [54] approached the problem from what is essentially the opposite (and arguably easier) direction: they investigated how CFT data such as fusion rules could be extracted from a function obtained as a conformal block. Recently, these ideas have been further developed

into the "pattern-of-zeros" approach [55], which has some similarities with the filtration method: the pattern of zeros S_a simply records the number of powers with which the wavefunction vanishes as a particles are brought together (in the later references, this is supplemented with other data from the CFT). Note that this method cannot accommodate the possibility of more than one fusion channel, and it was found in the references above that the mapping from symmetric polynomials to patterns of zeros is neither injective nor surjective: it is unknown what conditions need to be placed on the S_a to obtain a valid polynomial, and some values of S_a correspond to multiple Hall states.

A second approach has been the description of FQHE states via Jack polynomials [56], which are a basis for the space of symmetric polynomials, dependent on a free parameter α. In [57] it was first noticed that these polynomials obeyed a clustering property for certain values of α, and in [58] and elsewhere single elements of this basis were used to build clustered FQHE states. The Jack polynomial states also have a close relationship with the thin torus limit [59] which maps the two-dimensional many-body problem onto a one-dimensional lattice model. In this limit the Jack polynomial wavefunctions map on to their dominant (or "root") monomial term.

Recently it was conjectured [60] and then proved [61] that the (k, r)-admissible Jack polynomials are certain correlation functions of $\mathcal{W}A_{k-1}$ CFTs. While this is a large class of theories, the only unitary theories in this class are the Read-Rezayi series (defined in [62]) of parafermionic theories [60].

It is known that the above methods have a limited range of applicability and cannot uniquely describe all FQHE states. This was first pointed out in [63], in which a counterexample was constructed using the S_3 family of CFTs. The corresponding projection Hamiltonian requires the zero-energy wavefunctions to vanish as four powers of separation when four particles approach each other, but as was pointed out in [64], there are two linearly-independent ways in which that can happen. None of the approaches above is (currently) capable of providing this extra information.

In section 4.3.5 I employ another approach, the "filtration method," which was introduced by Read [65, 66] in the context of counting excitations of the parafermion states [62]. This method consists in producing a filtration over the space of symmetric polynomials with a clustering prop-

erty by observing when the wavefunction vanishes as certain numbers of clusters of various sizes are formed, leading to a maximally clustered form called a "residue." From the set of residues, one may construct a basis for the excitations.

1.2.3 Outline of new results

Because chapter 4 describes work in progress, the ratio of background material to new results is accordingly higher. I now summarize the new results presented in this part of the thesis.

In section 4.3.3 I introduce a new basis of translationally-invariant symmetric polynomials (4.71) (and additionally (4.76)) which is more readily applicable to filtration arguments than the polynomial basis proposed in [64] (equations (4.62) and (4.63)). I prove that the polynomials in this basis are linearly independent.

Stable quantum Hall phases are described by gapped Hamiltonians and unitary CFTs [53], so in sections 4.3.5 and 4.3.6 I attempt to modify the set of basis functions (4.97) found by Read in [67] for the generic SCFT in order to describe specific minimal SCFTs. These modifications correspond to the inclusion of extra projection terms in the Hamiltonian. Section 4.3.5 deals with the $SM(2,8)$ model, building on related results of Feigin, Jimbo and Miwa [68] for the $M(3,p)$ series of Virasoro minimal models; here I use the correspondence $SM(2,8) = M(3,8)$. This model is non-unitary, but it has the advantage of being simple to deal with: I show how a basis of excited states (4.109) may be obtained from the generic superconformal basis by adding extra couplings between clusters of particles. I confirm that, with this change, the state counting correctly reproduces a fermionic character for $SM(2,8)$, (4.112), meaning that all singular vectors have been dealt with. I discuss the form of the extra term which must be added to the projection Hamiltonian in order to allow only these states.

In section 4.3.6 I attempt to repeat the procedure on a unitary SCFT, the $SM(3,5)$ model (which describes the tricritical Ising model [69]), and show that existing characters for this model cannot be reproduced with this procedure. This strongly implies that the filtration argument must be generalized in order to deal with this state. Similar difficulties are encountered with the $SM(2,4k)$ series of models for $k \geq 3$, implying that this behavior is the rule rather than the exception.

In this section I also describe an alternative approach to the tricritical Ising model: rather than begin with the basis functions, I use an overcomplete basis of states given by the action of the Virasoro modes and modes of a chiral boson and determine the dimension of the space spanned by these states at each level. This allows me to obtain the characters of the tricritical Ising model at finite particle number (in equation (4.136)), which provide clues to which additional states must be removed by the projection Hamiltonian. I obtain the explicit form of the level three singular vector (the first state which must be removed) in the seven-particle system in equation (4.137), using the basis obtained for the generic superconformal theory.

1.2.4 Future directions

As mentioned above, this project remains a work in progress. Obviously the main goal is to understand what additional states must be removed by the projection Hamiltonian in order to obtain the tricritical Ising model or some other unitary SCFT. What results I have on this model (given in section 4.3.6) imply that the structure of states will look very different than those of a SCFT at generic c, and probably involve some formulation of clusters of clusters of particles.

Intermediate steps to this goal may be provided by studying other, nonunitary models with related but simpler structure. In section 4.4 I discuss one promising series of models, the $SM(2, 4k)$ SCFTs for $k \geq 2$; the $k = 2$ case is treated in section 4.3.5 and perhaps the useful features of its solution carry over to other models in the series.

An effective solution to this problem will require incorporating data from the desired CFT into the polynomial basis used, in a more extensive way than is done in section 4.3.4. The clustering properties I make use of in the filtration argument are obviously related to operator-product expansions of the particle operators in the underlying CFT, but it may be necessary to know this expansion to very high order in order to capture the complete singular vector structure.

Chapter 2

Minimum spanning trees: Bethe lattice

In this chapter, I am concerned with a treatment of the minimum spanning tree (MST) problem in the mean-field approximation. This problem asks us, given a connected graph with a real cost on each edge, to find the spanning tree which minimizes the total cost of the occupied edges. In order to obtain an interesting statistical mechanics problem, I fix the geometric structure of the graph and let the edge costs be quenched independent random variables, thereby obtaining a random MST ensemble. I am able to analyze this model due to a close relation between one of the simplest algorithms for solving the MST problem [8] and bond percolation. My principal motivation for studying the random MST problem comes from a strongly disordered spin glass model proposed by Newman and Stein ([23] and section 2.1.5): The procedure for generating the ground states of this model corresponds exactly to the MST problem I study. These concepts and arguments are laid out in an appropriate level of detail in section 2.1.

In section 2.2 I formulate the random MST problem on the Bethe lattice with appropriately chosen wired boundary conditions. The main objects I will be concerned with are D_{p}, the fractal dimension of paths on the MST, and D_{MSF}, the fractal dimension of connected components of the MST visible within a window. This latter quantity is important because it sets the upper critical dimension d_c for my model, above which mean-field theory is exact. I find (in section 2.2.7) $D_{\mathrm{p}} = 2$ and $D_{\mathrm{MSF}} = d_c = 6$. The latter result differs from the value $d_c = 8$ suggested in [23] via a scaling argument; in section 2.1.6 I show where the assumptions of that argument are in error.

My result implies that the strongly-disordered spin glass has infinitely many ground states above $d = 6$ but only one pair below this value. In section 2.3 I argue my results also apply to random ensembles of graphs which have a locally treelike structure.

In the following chapter I describe how corrections to the mean-field results may be computed as an asymptotic series in $\varepsilon = d_c - d$. Because the behavior of D_{MSF} is trivial below d_c (there is only of order one space-filling connected component), there I concentrate on the fractal dimension of paths on (portions of) the MST.

2.1 Minimum spanning trees, percolation and spin glasses

In this introductory section, I define and motivate the models and concepts of statistical mechanics that I will make extensive use of in this chapter and the next. These considerations form the foundation for the calculations I set out to perform.

I begin in section 2.1.1 by defining my object of study, the minimum spanning tree of a graph, and deducing several properties that permit its efficient solution. I summarize two of the simplest algorithms for computing the MST. In section 2.1.2 I turn to the subject of bond percolation, in particular its description as a geometric critical phenomenon and its scaling behavior. I discuss a geometric criterion for identifying the breakdown of mean-field theory, and introduce the related model of invasion percolation. In section 2.1.3 I link the MST problem and percolation — this correspondence is the key principle which enables all the computations in this chapter and the next.

In section 2.1.4 I give a quick introduction to spin glass models in general. This is followed in section 2.1.5 by a more detailed discussion of a strongly-disordered, exactly soluble spin glass model proposed by Newman and Stein in [23]. This model is of interest because the problem of finding its ground states may be mapped onto the MST problem; I describe this mapping in detail. My discussion of the Newman-Stein model continues in section 2.1.6, where I discuss its ground state structure. I also give a non-rigorous scaling argument establishing the upper critical dimension, which is validated by the qualitative calculations in the next section.

2.1.1 Minimum spanning trees (MSTs): properties and algorithms

Definitions

An instance of the minimum spanning tree (MST) problem is defined by an undirected, connected graph G with vertex set V and edge set E with an associated cost ℓ_e assigned to each edge $e \in E$. A solution of this problem — a minimal spanning tree — is an edge subset $T \in E$ which minimizes the sum of the costs of the included edges,

$$\ell(T) = \sum_{e \in T} \ell_e, \tag{2.1}$$

connects all vertices (hence, "spanning") and contains no closed cycles of edges, which makes it a tree. If the costs ℓ_e are positive, then any spanning subset of the edges which has minimum cost is automatically a tree, and if the costs are distinct the minimum spanning tree is guaranteed to be unique. Specifically, I will take the edge costs to be independent and identically distributed (iid) according to some continuous distribution, which ensures the uniqueness of the MST with probability one.

In what follows, I will view the cost (2.1) as an energy and the MST problem as a problem of finding the ground state of a classical system whose possible configurations are spanning trees. I will take the graph G to be fixed and the edge costs to be random variables, which makes the MST equivalent to a classical system at zero temperature with quenched disorder: the energy must be minimized for a fixed set of edge costs before any averaging over realizations of the edge costs is performed. The annealed average corresponds to a statistical ensemble in which each spanning tree of G is equally likely, which I refer to as uniform spanning trees (USTs).

In the following discussion, it will be useful to generalize this definition to a spanning forest, which is a set of connected subsets of E and a disjoint partition of V, such that the vertices in each part are spanned by an edge subset which has no cycles and no edges connect different parts. Note that the vertex set partition may include isolated vertices. If a spanning forest is a subset of the MST, I refer to it as a minimal spanning forest (MSF).

Properties of the MST

For a given realization of edge costs, the MST enjoys many convenient properties which make possible simple and efficient algorithms for the solution of the MST problem.

A reparameterization property of the MST comes from considering the converse of the MST problem, which will often be the natural point of view in what follows. With G fixed, the edge costs determine the MST, but only a limited set of information about the edge costs may be deduced from the MST. Obviously this doesn't include the numerical value of any ℓ_e: the MST for two realizations of edge costs that have the same relative order $\ell_{e_1} < \ell_{e_2} < \ell_{e_3} < \cdots$ will be the same, as pointed out in [10] and elsewhere. In what follows, I will make use of this invariance and without loss of generality take the distribution of edge costs to be uniform on the unit interval. I denote these rescaled costs by p.

In fact, the MST defines only a partial ordering on the set of edge costs, as evidenced by the fact that the most efficient known deterministic algorithm for computing the MST on an arbitrary graph [70] has significantly faster asymptotic running time than sorting all the edges of that graph by cost. We may generate the set of inequalities among the edge costs from a given MST by repeated application of any of the three geometric properties given below.

On a connected graph each spanning tree contains $|V| - 1$ edges, so if there were a bijection between spanning trees and orderings of costs (or if each tree arose from the same number of orderings), each tree would be equally likely to be minimal. This would imply that the statistical ensemble of MSTs would be identical to the ensemble of USTs [47, 71], which is known not to be the case [18, 21, 33, 42].

Further well-known geometric properties of the MST, proved in [5], are

1. The strong cycle property: for any edge not on the MST, there exists a unique closed cycle of edges containing that edge such that all other edges on the cycle belong to the MST and are less expensive than the given edge;

2. The strong cut property states that for any edge on the MST, there exists a unique disjoint partition of the set of vertices into two parts connected by that edge such that all other edges

joining vertices between the two parts are more expensive;

3. The minimax path property states that the path on the MST connecting any two vertices minimizes the cost of the most expensive edge encountered along that path. I refer to this most expensive edge (which is unique if edge costs are distinct) as the barrier between those two vertices. In fact, this statement may be made more precise: paths on the MST may be uniquely characterized in terms of a geodesic property, which is defined and proved in section 3.2.1 (we will not need it in this chapter).

Any two of these statements implies the third.

A further monotonicity (or "greediness") property is especially relevant for the operation of MST algorithms [4]. Assume we are given some spanning forest F on G, which is not necessarily minimal. There exists some spanning tree T_F of G which minimizes the edge cost subject to the constraint that it contain all edges of F. If we select a connected component of F and find the edge e of minimum cost incident on only one vertex of that component, that edge e must be in T_F.

This follows readily from the strong cut property: we may construct a derived graph G' by contracting all connected components of F to single vertices. Spanning trees of G' are in one-to-one correspondence with spanning trees of G containing F. The statement follows from considering the cut where the selected component (vertex) is one part of the cut, and all other vertices of G' are the other part.

MST Algorithms

Adding e to F yields a forest with one fewer component, and so this process can be iterated to construct T_F. Hence the term "greedy": at each step we make a locally optimal decision (the addition of e) and we never need to revise past decisions. In particular, if we begin with the empty forest consisting of isolated vertices and no edges, $G = G'$ and we construct the MST of G. The only detail remaining to be specified is a criterion for selecting components of F, which leads to the two main simple algorithms for computing the MST:

1. The Jarník-Prim-Dijkstra algorithm [9] (conventionally referred to as Prim's) results from

selecting the same component at each step. As the algorithm runs, we grow a single large connected component or "cluster" starting from the initial vertex.

2. Kruskal's algorithm [8] results from selecting the component with cheapest bordering edge e remaining: in other words, at each step one adds the cheapest edge not yet occupied that does not create a cycle when added to the set of occupied edges. Generically, at intermediate stages of the algorithm several clusters are present.

In either case, the algorithm stops when a spanning tree is formed after $|V|-1$ edges are occupied, and from the greedy property this spanning tree is the MST. In practice both algorithms require additional procedures for dealing with identical edge costs, which are not of interest here. The performance of either algorithm can be improved through the judicious choice of data structures.

2.1.2 Bond and invasion percolation

Bond percolation

Percolation has been extensively studied by both physicists and mathematicians; I refer the reader to the reviews [39, 72] for further information.

Bond percolation on a graph G is defined by taking each edge to be independently occupied with probability p or unoccupied with probability $1-p$. We then study various geometric properties of the clusters formed by the occupied edges. The most basic of these, for graphs that are a portion of a regular lattice, is the percolation threshold: for low p any given site (which we can take to be the origin without loss of generality) will be in a cluster of finite extent. For large p the origin is likely to be in a cluster of "infinite" extent, which we may define more carefully by working with a system of large, finite size, considering clusters which are of extensive size, and taking the limit as the system size goes to infinity. An alternative definition of "infinite cluster" which I will use below is to consider those clusters extending from a neighborhood of the origin to the boundary of the system.

Remarkably, these two regimes of behavior (connected and not connected to a distant boundary) are separated by a critical value of p which becomes a sharp threshold in the limit of infinite

system size limit. Specifically, there exists a percolation threshold p_c below which the origin is in a finite cluster with probability strictly one and above which there is some nonzero probability to be in an infinite (spanning) cluster. This shows that percolation is an example of a geometric critical phenomenon, in which scaling behavior is obtained in the absence of thermal fluctuations — here p is analogous to temperature, or (more accurately) the relevant perturbation which takes us across the critical fixed point at p_c. In fact, analogues of all the standard critical exponents of thermodynamic phase transitions may be defined in a way that I now briefly summarize.

Scaling properties of percolation

I now present some of the scaling properties for percolation on an infinite, finite-dimensional Euclidean lattice [39]. The percolation threshold p_c is non-universal: its value depends on the microscopic structure of the lattice or class of graphs considered. For hypercubic graphs in dimension d (*i.e.* $\mathbb{Z}^d$, with edges connecting nearest neighbors), p_c lies strictly between 0 and 1 for $d > 1$, while for $d = 1$, $p_c = 1$.

For $p < p_c$, there are finite clusters only. As p approaches p_c from below, the typical length scale ξ of a cluster (the correlation length) diverges as $\xi \sim (p_c - p)^{-\nu_{\text{perc}}}$. For $p > p_c$, with probability one, there is a single infinite cluster and, for $p < 1$, a non-zero density of finite clusters not connected to infinity. The correlation length ξ is now defined by the size of finite clusters, and diverges with the same power law $\xi \sim (p - p_c)^{-\nu_{\text{perc}}}$ as p_c is approached from above. The fact that the percolation transition is governed by a single diverging length scale ξ is the essential reason why it may be described in the standard language of thermodynamic phase transitions.

The critical exponent β_{perc} is defined via the probability that the origin is on an infinite cluster above p_c: we have $P_\infty(p) \propto (p - p_c)^{\beta_{\text{perc}}}$. The exponents α_{perc} and γ_{perc} are less intuitive and are related to the cluster size distribution: the total number of clusters has a singular part which diverges as $|p-p_c|^{2-\alpha_{\text{perc}}}$, and the mean number of vertices on a finite cluster diverges as $|p-p_c|^{\gamma_{\text{perc}}}$.

The exponents defined above take the same values on either side of the threshold and obey the standard scaling relations, such as the Rushbrooke equality $\alpha + 2\beta + \gamma = 2$. Below an upper critical dimension d_c, additional hyperscaling relations hold, namely $d\nu = \gamma + 2\beta$, where d is

the dimensionality of space. For $d > d_c$ the exponents are independent of d and take on their mean-field values, which are $\alpha_{\text{perc}} = -1, \beta_{\text{perc}} = 1, \gamma_{\text{perc}} = 1$ and $\nu_{\text{perc}} = 1/2$. The upper critical dimension for percolation is $d_c = 6$, and exactly at $d = d_c$ the power laws given above receive logarithmic corrections which I do not consider here.

Exactly at $p = p_c$, there is a power law distribution of cluster sizes up to infinite size. There is strictly zero probability for any site to be on an infinite cluster, but we may define the notion of an incipient infinite cluster by looking at the largest (extensive) finite clusters in the system. These clusters have fractal dimension $D_{\text{perc}} = d - \beta_{\text{perc}}/\nu_{\text{perc}}$ for $d < d_c$ and $D_{\text{perc}} = 4$ above d_c. This analysis holds in the vicinity of p_c for length scales less than ξ; for $p > p_c$ on scales $L \gg \xi$ the clusters have $D_{\text{perc}} = d$. Equating these relations at scales $L \simeq \xi$ yields the hyperscaling relation for D_{perc} given above. For $p < p_c$ and scales $L \gg \xi$, $D_{\text{perc}} = 0$ and the large clusters grow only logarithmically with the system size.

Cluster proliferation and the breakdown of mean-field theory

In this chapter and the next, I take pains to formulate my arguments in terms of geometric concepts such as cluster size and connectivity. I should therefore point out how the breakdown of hyperscaling and mean-field theory above and below the upper critical dimension may also be phrased in geometric terms [39, 46].

Essentially, what happens is that for $d < d_c$ there is only one (or rather, of order one) incipient infinite cluster, while above d_c there are extensively many. More rigorously, we can define an additional proliferation exponent $\#_{\text{perc}}$ so that for $d \geq 6$, there are of order $L^{\#_{\text{perc}}}$ large such clusters intersecting a box of size L, where $\#_{\text{perc}} = d - 6$, while $\#_{\text{perc}} = 0$ for $d \leq 6$. Factors of $\xi^{\#_{\text{perc}}}$ will then enter all the scaling relations, since we need to include contributions from all spanning clusters, and this d-dependence will cancel the d-dependence in hyperscaling relations, fixing the critical exponents at their mean-field values.

In this way, we can identify the upper critical dimension from purely geometric considerations, rather than from the conventional analysis of the relevance of fluctuations induced by low dimensionality. This will be useful for my purposes since I am dealing with geometric critical phenomena

in the absence of thermal fluctuations.

"Dynamic" bond percolation

We may view percolation as a simple example of a quenched random critical phenomenon: consider a graph with iid edge costs as described above. If $P(\ell_e)$ is the distribution of any edge cost, then if we occupy all edges with cost less than or equal to some ℓ_0 we obtain a realization of bond percolation with parameter $p = p_0$ given by

$$p_0 = \int_{-\infty}^{\ell_0} d\ell_e \, P(\ell_e). \tag{2.2}$$

Without loss of generality, we may rescale all the edge costs according to (2.2) and take the distribution to be uniform on the unit interval, since all that matters in this process is whether a cost is greater or less than p_0.

We can then obtain a dynamic percolation process by raising p_0 from 0 to 1. This grows a set of clusters starting from the configuration consisting of all isolated vertices by adding edges in order of their quenched cost, and terminates when all edges are occupied. This definition may appear trivial, but it merits discussion due to its connection with Kruskal's algorithm (below).

Invasion percolation

An alternative dynamic model of percolation is invasion percolation, which was defined in [73] to more closely model the percolative behavior of fluids in real materials (see, also, *e.g.*, [39]). Here we take a graph with quenched edge costs as before, but grow a single cluster starting with a given initial vertex (meant to represent a site of fluid injection). At each step, we occupy the least expensive edge on the boundary of the cluster (*i.e.*, which is incident on at least one vertex in the cluster).

Note that, in this process, the costs of successively occupied edges may not increase monotonically: new border edges become available as the cluster grows and these may be cheaper than previously occupied edges. In fact, in an infinite system as the process is run for a long time the newly added edges will all have costs $\leq p_c$, where p_c is the bond percolation threshold on the same

graph [73, 74]. In this long-time limit properties of the invasion percolation cluster approach those of a critical percolating cluster, so its fractal dimension is $D_{\text{inv}} = D_{\text{perc}} = 4$ for $d \geq 6$ [73, 74].

2.1.3 MST algorithms and percolation

Remarkably, the percolative processes mentioned above enjoy a close relationship with the Kruskal and Prim algorithms (respectively), which is essential for all calculations done in this chapter. In both cases, we obtain an algorithm for generating the MST from the percolative process by imposing an additional "no-cycle" condition: we consider edges in the order dictated by the percolation process, but if occupying an edge would form a cycle with other previously occupied edges, we leave that edge unoccupied and move on to the next. Note that this condition, if thought of as an interaction, is highly nonlocal both in space (we may need to examine edges arbitrarily far away to determine whether they form a cycle) and in time (we may need to examine edges which were accepted arbitrarily long ago).

Kruskal's algorithm and bond percolation

Inspection of the respective definitions shows that Kruskal's algorithm is simply the dynamic bond percolation process with the no-cycle condition imposed [75]. Note that that only the order in which edges are considered matters (as a consequence of the reparametrization property of the MST [10]). The process is fully characterized by the variable p_0, and distinct probability densities $P(\ell_e)$ give rise to the same probability measure on MSTs. Without loss of generality, in what follows I will take ℓ_e to be uniformly distributed between 0 and 1, so $p_0 = \ell_0$, as above.

At any intermediate stage of the algorithm (with "time" being measured by the number of edges that have been considered, or equivalently the value of the percolation parameter p_0) the set of occupied MST edges is a subset of the edges belonging to the percolation clusters obtained at p_0. This defines a minimal spanning forest which I denote MSF(p_0); MSF(1) is the full MST. The object MSF(p_0) for $p_0 < 1$ will also be of interest to us, particularly in chapter 3; for this reason I will refer to the growth of MSF(p_0) as p_0 is raised as the "Kruskal process" in what follows below.

Prim's algorithm and invasion percolation

The connection between Prim's algorithm and invasion percolation has been recognized a number of times [10, 11, 17, 23]. There is a complication in that the algorithm generates the full MST on a finite graph, but for sufficiently high dimensions in an infinite system the tree obtained through this process will not be spanning, even as the number N of edges considered goes to infinity. Because this tree is dependent on the choice of initial vertex $\mathbf{x}$, I denote it by $\lim_{N\to\infty} T_N(\mathbf{x}) = T_\infty(\mathbf{x})$. Clearly, the set of vertices on the invasion cluster and on the invasion tree produced by Prim's algorithm are the same, so $T_\infty(\mathbf{x})$ will have the same fractal dimension as the infinite percolating cluster at p_c.

2.1.4 Spin glasses

The term "spin glass" can refer to any of a number of models of disordered magnetic materials. A vast literature exists on the subject; I direct the reader to the reviews [25, 76, 77] for further background than is provided here.

The prototype for such materials is a non-magnetic metal doped with magnetic atoms which interact via the RKKY interaction, whose sign oscillates as a function of distance. Because the structure of the resulting alloy is amorphous, a generic magnetic moment will have both ferromagnetic and antiferromagnetic interactions of varying strengths with its neighbors.

Edwards and Anderson [78] argued that we must extend our notion of ferromagnetic order in order to accommodate this pattern of interactions: such a system will always have a vanishing net magnetization $m = \frac{1}{N}\sum_i \langle s_i \rangle$ but at sufficiently low temperatures will have a nonzero value of $q_{\mathrm{EA}} = \frac{1}{N}\sum_i \langle s_i \rangle^2$, reflecting the "freezing in" of the spins to definite, fixed but spatially random orientations. This leads to the two essential qualities for spin glass behavior:

1. Quenched disorder of coupling strengths, which breaks translational invariance of the system. We recall that quenched averages of thermodynamic quantities must be computed using a fixed sample of couplings, which are only averaged over in the end.

2. Frustration, in which the product of signs around a cycle of couplings is negative. This means

there does not exist a ground state in which all couplings are satisfied, *i.e.* in which the spin orientation minimizes the energy of all couplings individually.

The simultaneous presence of these two phenomena produce a complicated free energy landscape in the system's configuration space, which (unlike a non-random system) may have (extensively) many local minima. This roughness is responsible for the experimentally observed consequences of glassy behavior, such as extremely long equilibration times. It also makes analytic or numerical treatment of spin glass models extremely difficult with existing methods.

The above features are represented in more abstract form in the Edwards-Anderson (EA) Ising spin glass model [78], defined by the Hamiltonian

$$\mathcal{H} = -\sum_{i<j} J_{ij} s_i s_j \tag{2.3}$$

for Ising spins $s_i = \pm 1$, where J_{ij} are quenched random variables taking both positive and negative values. The underlying graph is typically taken to be a d-dimensional hypercubic lattice, and the couplings are taken to be among nearest neighbors and iid according to some lattice-independent distribution, for example a Gaussian with zero mean and fixed width.

The ground state of an EA model on a planar graph may be found in polynomial time via a mapping to a network flow problem, but for $d > 2$ the problem of determining the ground state is computationally costly. In fact, for this general (nonplanar) case the related problem of determining whether the ground state energy is less than some bound is NP-complete [79], and hence is believed to be impossible to solve in polynomial time. For a lengthier discussion of NP-completeness and related issues from computer science, see [80].

Controversial aspects

The mean-field solution of the Edwards-Anderson model is formulated on the complete graph [81], in which there is an (appropriately scaled) coupling between any pair of spins. The solution of this model was first considered by Thouless, Anderson and Palmer [82] and fully achieved by by Parisi and collaborators [83], under the name of replica symmetry breaking. This solution displays

a number of remarkable features, most notably the presence of an infinite number of ground states, not related by any symmetry transformation, described by an order parameter *function*.

It is generally acknowledged that the Parisi solution of the mean-field theory is correct, and progress has been made on establishing the solution on a rigorous mathematical footing. What remains controversial is what connection, if any, the mean-field solution has with the EA model in finite dimension with short-range interactions; *i.e.* whether the Parisi solution constitutes a sensible starting point for perturbation theory. Because of this uncertainty, even such basic features as the lower and upper critical dimensions of the EA model and the number of ground states in the spin-glass phase (finite or infinite) have not been conclusively determined [77].

Perhaps the most well-developed alternative scenario to replica symmetry breaking is the droplet scaling model [84], which contradicts the former picture in several ways. The droplet scaling model is a phenomenological picture in which the spin glass only has two ground states, related by a global spin flip symmetry. Thermal fluctuations manifest themselves in the form of clusters of flipped spins, with the size distribution of these clusters falling off as a power law.

2.1.5 Strongly-disordered spin glass model and MSTs

Newman and Stein [23] (see also [11]) defined a strongly-disordered spin glass model, constructed so as to be exactly soluble. The salient feature of this model is the presence of extensively many ground states above an upper critical dimension, with only a single pair below this dimensionality, thus providing a not implausible scenario for the ground state multiplicity questions raised in the previous section. The problem of finding the ground state of this model may be mapped onto the MST problem, which provides one of my main motivations for studying this system.

Definition

The Newman-Stein (NS) model may be considered as the strong disorder limit of the EA model on an arbitrary lattice in Euclidean space: the couplings J_{ij} between the Ising spins are still iid, but now we take the width of their distribution to be extremely large. In fact, as it turns out, this width will need to grow extensively with the system size, for reasons described below. We take

the coupling distribution to be symmetric, so that the sign of each J_{ij} is positive or negative with probability 1/2, independent of the coupling magnitude $K_{ij} = |J_{ij}|$.

The other substantial difference between the NS model and the EA model is that Newman and Stein explicitly specify the boundary conditions to be used: each spin on the boundary of the lattice is fixed to be either ± 1 arbitrarily, chosen independently of any couplings or spin values in the bulk. These are referred to as wired boundary conditions, and their importance will become clear in what follows. (Note that I have already referred to wired boundary conditions for MSTs in section 1.1.3: these are the spin glass wired boundary conditions under the mapping to the MST problem, as will be discussed further later on.)

Ground state algorithm

The primary utility of going to the strong disorder limit of the coupling distribution is that we may assume that for any subset of edges, the coupling with largest magnitude dominates not only the magnitude of all the other couplings but also their sum. This condition may be met if each K is larger than at least twice the next smallest K, which requires the width of the distribution to grow with the number of couplings in the system — hence this model is not thermodynamically realistic.

Once this condition is met, the ground state of the model for a given choice of wired boundary condition may be found by a greedy procedure. In Newman and Stein's paper, they used a procedure related to Prim's algorithm as follows. In order to find the ground state orientation (either ± 1) of the spin at site i, we start by finding the largest K_{ij} among all edges incident on i. By the strong disorder assumption, this coupling must be satisfied, so this fixes the relative orientation of spins s_i and s_j.

We then repeat the process, finding the next largest K' incident on this cluster of two sites, which will fix the orientation of the cluster relative to some third spin s_k, and so forth. We do not need to consider edges that would form cycles in the cluster, since by the strong disorder assumption the corresponding coupling is too weak to change any orientations that have previously been determined — *i.e.*, frustration has no effect on the ground-state structure of this model, which is the key feature enabling this model's exact solution. The process terminates when the cluster

grown from i (whose spins have been determined up to a global flip) touches the wired boundary, which fixes the absolute orientation of all spins on the cluster. The orientation of remaining spins may be determined by starting over and repeating the procedure for some site not on the cluster, growing the tree until it intersects some previously fixed spin (either on the boundary or on another tree), and continuing until the ground-state orientation of all spins has been fixed.

Connection with MSTs

This algorithm may be seen to compute the MST with the following identifications. First, the spin glass couplings are considered in descending order of magnitude, so the ordering must be reversed: the progress of the spin glass algorithm is identical to that of Prim's algorithm on the same graph with edge costs $\{-K_{ij}\}$, up to when the first tree hits the boundary. In order to complete the identification, we take all boundary spins to be connected by edges of less cost (greater magnitude) than any edges in the bulk — this is the wired boundary condition for MSTs mentioned in section 1.1.3, now motivated by the Newman-Stein spin glass model. MSTs with wired boundary conditions were previously studied for their own sake in [7].

The wired boundary condition means that when the first tree hits the boundary, it then becomes connected to all boundary spins, and that the final collection of all the trees is a MST on the original graph plus the boundary edges. The algorithm originally presented by Newman and Stein is not identical to Prim's algorithm, because in order to generate the complete MST/ground state we must periodically restart the growth of a new tree from a new vertex in the interior, but as noted in section 2.1.1 there are many possible algorithms which all yield the MST.

Because the ultimate pattern of satisfied couplings may be mapped onto the MST, we may use any algorithm for finding the MST to compute it — in particular, we may use Kruskal's algorithm and its relation to dynamic bond percolation, discussed in section 2.1.3 . Once the MST is found, the spin orientations follow from the signs ϵ_{ij} of the couplings and the boundary spin orientation. The treelike nature of the set of satisfied couplings means that there is no frustration, and since the MST may be found in polynomial time, so may the ground state of this strongly disordered spin glass.

2.1.6 Ground states and critical dimension of the NS model

Newman and Stein constructed their model so that it would be exactly soluble by the above procedure. The model was intended to probe ground state multiplicity in spin glasses [23, 77]. In order to discuss their result, we must examine more rigorously how thermodynamic states are defined.

Constructing ground states of spin systems

Roughly speaking, we may think of a thermodynamic ground state as corresponding to one of the local minima in the free energy landscape introduced in section 2.1.4, which is separated from other local minima by barriers whose height is extensive with the system size.

In what follows we need only consider ground states, which can be identified (in an infinite system) as those spin configurations whose energy cannot be lowered by changing the values of any finite set of spins. A ground state on any finite portion of the lattice may be uniquely specified by a choice of wired boundary conditions. We construct the thermodynamic limit via a two-step process: we first let the size L of the system go to infinity by adding new spins and iid couplings to those already present. In order for thermodynamic limit to exist, we need to select boundary conditions for each of the sequence of finite-size systems so that the spin configuration (more specifically, any correlation function) visible within any subregion of the system — a "window," of size $W \ll L$ — converges to well-defined values as $L \to \infty$. We then let the size of the window go to infinity (only after having let $L \to \infty$ first) and require that the configurations agree where the windows overlap. This is the rigorous definition of what is meant by the thermodynamic limit of a lattice model. These precautions are necessary due to the uncertainty over the number of ground states of the EA spin glass, as discussed in section 2.1.4.

Ground state multiplicity of the NS model

As described above, Newman and Stein constructed the ground state of their model based on growing a number of invasion trees. The relative orientation of all spins on a connected component are fixed by the bulk couplings and the strong disorder assumption, but the absolute orientation

(or relative orientation of different components) is completely fixed by the orientation of the spin where the tree intersects the wired boundary. As pointed out above, this spin is unique, and if we change the wired boundary conditions so that its orientation is reversed, all spins on the connected component will change.

This means that the logarithm of the number of possible different ground states of a NS model in a box of size L is bounded by $L^{d-1}\log 2$. This is only a bound since some boundary spins may not be connected to any interior spins by the invasion tree algorithm, and changing them will not change the ground state. However, in light of the discussion above, the real quantity of interest is how many ground states are distinguishable in a window of size $W \ll L$ within that system as boundary conditions are changed. The logarithm of this number, $\log \mathcal{N}(W)$ is simply the number of connected components of the MSF of size extensive in W which intersect the window. There are potential ambiguities posed by the presence of subextensive components and components which intersect the window more than once (*i.e.,* which are connected somewhere outside the window), but I assume these effects are not significant.

Therefore, we have

$$\ln \mathcal{N}(W) = N_{\text{MSF}}(W)\log 2, \tag{2.4}$$

where N_{MSF} is the number of large MSF components visible in the window, which I take to behave as $W^{\#_{\text{MSF}}}$. This establishes that the upper critical dimension of the NS model is the same as that of the MST problem, defined geometrically via spanning cluster proliferation as described in section 2.1.2. Above d_c, $\#_{\text{MSF}}$ is nonzero, N_{MSF} will scale with W, there will be extensively many MSF components visible in the window and the NS model will have extensively many ground states, not related by any global symmetry. Below d_c, $\#_{\text{MSF}} = 0$, N_{MSF} will be of order one for any W, and the NS model will only have one (or finitely many) ground state pairs, which are related by a global spin flip.

This connection to the NS model is my main motivation for studying N_{MSF}. In practice, since the MSF components are space-filling, we can arrive at this number by computing the fractal dimension of a single component, which I will do by integrating the probability that two points are on the same connected component.

Scaling argument for the critical dimension of the NS model

I now give a non-rigorous scaling argument determining d_c for the MST problem and NS models, which will be confirmed by the calculations of the next section. Recall that in Newman and Stein's ground state algorithm, we needed to restart the invasion tree process from new spins in the interior until the full ground state configuration was determined. In order to determine whether two well-separated vertices $\mathbf{x}$, $\mathbf{y}$ are on the same MSF component, we can first grow the full invasion tree from $\mathbf{x}$ to the boundary at infinity (yielding $T_\infty(\mathbf{x})$), then restart the process from $\mathbf{y} \notin T_\infty(\mathbf{x})$, stopping when the invasion tree hits either the boundary or $T_\infty(\mathbf{x})$. As noted in section 2.1.2, $T_\infty(\mathbf{x})$ is a fractal object with dimension $D_{\text{inv}} = 4$ for $d \geq 6$.

Therefore, as Newman and Stein argued, when $d > 2D_{\text{inv}} = 8$, there is a non-zero probability that $T(\mathbf{y})$ will miss $T_\infty(\mathbf{x})$ completely and connect to the boundary first [23]. This implies that $\mathbf{x}$ and $\mathbf{y}$ are on distinct components, so we have cluster proliferation above $d = 8$, hence $d_c \leq 8$.

This argument would establish $d_c = 8$ if the two fractal objects $T_\infty(\mathbf{x})$ and $T(\mathbf{y})$ intersected with probability of order one for $d < 8$, which would follow if they were statistically independent by the simple fact that $2D_{\text{inv}} < 8$. Since we only examine edges on the perimeter of an invasion tree while it is growing, $T_\infty(\mathbf{x})$ and $T(\mathbf{y})$ are independent when they are far apart. (Recall that the perimeter of an invasion tree is the set of all edges incident on only one vertex of the tree).

This assumption of independence breaks down when the trees are about to intersect, however: for this event to occur, we must add an edge to $T(\mathbf{y})$ which is on the perimeter of $T_\infty(\mathbf{x})$ and hence has already been examined. From our discussion of invasion percolation, we know that the costs of edges on the perimeter of $T_\infty(\mathbf{x})$ are not distributed like those of the edges in the bulk and, in fact, have cost strictly greater than p_c. This creates an effective surface repulsion between the two trees which lowers the critical dimension.

This argument can be expressed in terms of correlation functions as follows. The probability $G(\mathbf{z} - \mathbf{x})$ that a point $\mathbf{z}$ will be on $T_\infty(\mathbf{x})$ should behave as

$$G(\mathbf{z} - \mathbf{x}) \sim |\mathbf{z} - \mathbf{x}|^{D_{\text{inv}} - d} \tag{2.5}$$

for large $|\mathbf{z}-\mathbf{x}|$ [23, 39]. If $T_\infty(\mathbf{x})$ and $T(\mathbf{y})$ were noninteracting and statistically independent, the probability they would intersect would be proportional to

$$\int d^d\mathbf{z}\, G(\mathbf{z}-\mathbf{x})G(\mathbf{z}-\mathbf{y}) \sim |\mathbf{x}-\mathbf{y}|^{2D_{\text{inv}}-d}. \tag{2.6}$$

For $d < 2D_{\text{inv}}(d) = 8$ this is nonvanishing for large separations. This would establish the converse result of Newman and Stein and also mean that MSF components have dimension at least $2D_{\text{inv}} = 8$ for $d \geq 8$ (though Newman and Stein seem to have believed that this dimension would be 4). However, this estimate must be corrected to take the repulsive interaction into account.

In order to do this, we can make a scaling argument involving the properties of $T(\mathbf{y})$. Let ξ be the scale of the linear size of the tree at some stage of its growth; then in order for $T(\mathbf{y})$ to intersect $T_\infty(\mathbf{x})$ we need to take ξ of order $|\mathbf{x}-\mathbf{y}|$. If $T(\mathbf{y})$ were infinite, its perimeter would be infinitely large and it would never accept any edges more expensive than p_c, but because of its finite size there are only finitely many edges to choose from, so occasionally $T(\mathbf{y})$ will have to add an edge of cost slightly greater than p_c. Specifically we expect the size of the deviation $p - p_c$ to be of order $\xi^{-1/\nu_{\text{perc}}}$ for an invasion tree of size ξ; the probability that a perimeter edge of cost p will be accepted when $T(\mathbf{y})$ has a size of order ξ will be some function f_1 of the scaling variable $(p-p_c)\xi^{-1/\nu_{\text{perc}}}$. This function should have the limiting behavior $f_1(x \to -\infty) \to 1$ and $f_1(x \to \infty) \to 0$. Likewise, the costs of perimeter edges on $T_\infty(\mathbf{x})$ a distance of order ξ from $\mathbf{x}$ are distributed as $f_2[(p-p_c)\xi^{-1/\nu_{\text{perc}}}]$, for some other scaling function f_2 satisfying $f_2(x<0)=0$, $f_2(x\to\infty)\to 1$. The intersection probability is then proportional to

$$\int dp\, f_1[(p-p_c)\xi^{-1/\nu_{\text{perc}}}] f_2[(p-p_c)\xi^{-1/\nu_{\text{perc}}}] = \mathcal{O}\left(\xi^{-1/\nu_{\text{perc}}}\right). \tag{2.7}$$

This analysis covers the first opportunity for intersection; we assume subsequent intersection evens contribute a converging sum of terms which only modify the right-hand side of this probability by a constant factor. Multiplying this probability by the right-hand side of equation (2.6), we find the intersection probability should scale as $|\mathbf{x}-\mathbf{y}|^{2D_{\text{perc}}-d-1/\nu_{\text{perc}}}$. Referring to numerical estimates of

D_{perc} and ν_{perc} for $d < 6$ [39], the only consistent solution to

$$2D_{\text{perc}} - d - 1/\nu_{\text{perc}} = 0 \tag{2.8}$$

happens at $d = 6$; hence for the MST problem $d_c = 6$ and

$$D_{\text{MSF}} = \begin{cases} 6 & (d \geq 6), \\ d & (d \leq 6); \end{cases} \tag{2.9}$$

$$\#_{\text{MSF}} = \begin{cases} d-6 & (d \geq 6), \\ 0 & (d \leq 6). \end{cases} \tag{2.10}$$

In the next section, I obtain the same results through more rigorous calculations on the Bethe lattice. I mention that the authors of Ref. [10], who did not investigate D_{MSF} or $\#_{\text{MSF}}$, also stated that the critical dimension for MSTs is six, apparently because they were using Prim's algorithm and its connection to invasion percolation which has $d_c = 6$.

2.2 The MST problem on the wired Bethe lattice

In the previous section I defined the MST problem and identified quantities of interest (arising from, *e.g.*, the correspondence with the Newman-Stein (NS) spin glass model discussed in section 2.1.5); in this section I will get down to the business of actual computations on the Bethe lattice, which I will use as a mean-field approximation to the problem in Euclidean space. I begin in section 2.2.1 by describing the slightly unorthodox way in which I construct the mean-field approximation for this problem (again referring to the NS model, I adopt an analogue of their wired boundary conditions on the Bethe lattice). Keeping in mind the percolation-Kruskal correspondence discussed in section 2.1.3, in section 2.2.2 I review properties of the mean-field theory of percolation on the Bethe lattice, which I will then modify in section 2.2.3 by imposing the history-dependent no-cycles condition of Kruskal's algorithm. I obtain integral equations for all correlation functions of interest.

In section 2.2.4 I begin the analysis of the scaling limit of these correlation functions, by defin-

ing an appropriate Green's function and determining its scaling behavior. In section 2.2.5 I continue by discussing the scaling behavior of the two-point correlation function, and in section 2.2.6 I give arguments extending my analysis to general k-point correlation functions. Finally, in section 2.2.7 I review these scaling results and describe how mean-field fractal dimensions (in finite-dimensional Euclidean space) may be inferred from them.

2.2.1 Definitions

I develop a mean-field treatment of the MST problem by specializing to the case of the Cayley tree, which is a finite tree of degree $z = \sigma + 1$ at every vertex, except for those at the boundary. The boundary is taken to consist of all vertices a "radius" of M steps from the origin. Since the MST problem on a graph which itself is already a tree is trivial, I adopt wired boundary conditions in order to obtain an interesting problem. I do this by connecting all boundary vertices to each other with additional edges of infinitesimal cost, so that all boundary vertices are connected for any $p_0 > 0$ (see figure 2.1). This construction implements the wired boundary conditions of the NS spin glass model (see section 2.1.5).

The "thermodynamic limit" of this system is obtained by taking $M \to \infty$ and looking at properties of clusters a finite distance from the origin, far away from the boundary, as discussed in section 2.1.6. The result of this construction is referred to as the Bethe lattice; this is also the "weak limit" discussed in [48], and this defines the mean-field model, whose scaling properties will be identical to that of the MST on a Euclidean lattice of high dimension.

My use of wired boundary conditions means that when we run the Kruskal process, growing clusters can only touch the boundary once: two connections to the boundary would form a cycle. When the process is run to completion at $p_0 = 1$, we obtain the MST of the wired Cayley tree, which becomes a minimum spanning forest when we take the Bethe limit and neglect connections at the boundary (equivalently, examine the system in a window around the origin).

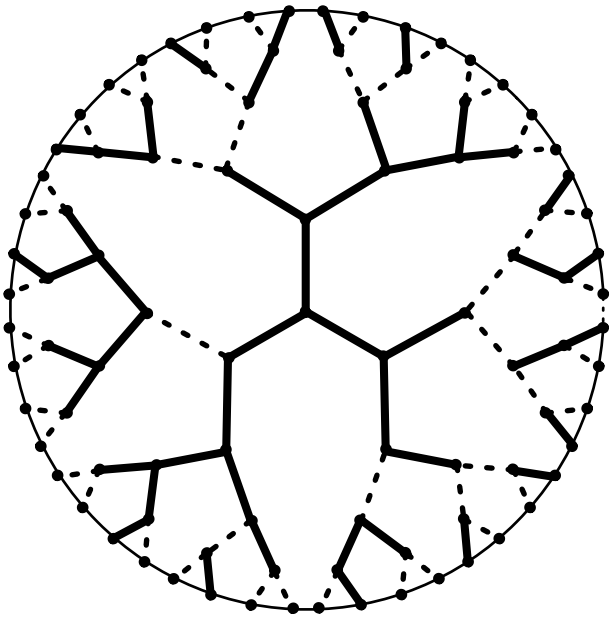

Figure 2.1: A sample realization of the MST (solid lines) on the wired Cayley tree. The graph shown has $\sigma = 2$ and radius $M = 5$. Vertices on the boundary are connected with boundary edges of vanishing cost, shown here as thin lines. All vertices in the interior are connected to the boundary only once.

2.2.2 Percolation on the Bethe lattice

As explained previously (section 2.1.3), I make progress by using the correspondence between the Kruskal process and dynamic bond percolation. Percolation on the Bethe lattice has been studied thoroughly [39, 85]. The order parameter in this case is the probability $P_\infty(p)$ that the origin (or any vertex in the Bethe lattice) is on a cluster connected to the boundary, as a function of the edge occupation probability. We obtain this by introducing the probability $F_m(p)$ that a given vertex m steps in from the boundary is *not* connected to the boundary through a given outward-facing edge. This is related to the percolation probability via $P_M(p) = 1 - F_M^{\sigma+1}(p)$ and satisfies a recurrence relation

$$F_{m+1}(p) = (1-p) + pF_m(p)^\sigma, \tag{2.11}$$

with $F_0(p) = 0$. The first and second terms on the right-hand side of (2.11) correspond to either the first outward-facing edge being unoccupied or that edge being occupied and no connection to the boundary being available through the σ branches at the next step.

In the Bethe lattice limit, all vertices a fixed distance from the origin are far from the boundary, and the probability of any of them not being connected to the boundary along a particular outward branch approaches a limit $\lim_{m\to\infty} F_m(p) = F(p)$, which is given by a stable fixed point of

$$F(p) = (1-p) + pF(p)^\sigma. \tag{2.12}$$

This has the trivial solution $F = 1$, whose stability is given by linearizing the recurrence about this solution for small deviations. The eigenvalue of the linearized recurrence for $F_m - 1$ is σp. For $p < 1/\sigma$, the solution $F = 1$ is stable and, with probability one as $M \to \infty$, no vertex is connected to the boundary. For $p > 1/\sigma$, the solution $F = 1$ is unstable, and the stable solution corresponds to a different root of (2.12) with $0 < F(p) < 1$; therefore, the percolation threshold on the Bethe lattice is $p_c = 1/\sigma$ and the percolation probability $P_\infty(p) = 1 - F^{\sigma+1}(p)$ is zero below p_c. Existence, uniqueness, and stability of these fixed points is proved in, for example, [86] (in different notation).

For $\sigma = 2$, (2.12) may be solved explicitly:

$$F(p) = \begin{cases} 1 & p \le p_c = \frac{1}{2} \\ \frac{1-p}{p} & p > \frac{1}{2}. \end{cases} \tag{2.13}$$

For general σ, we can expand in powers of $p - p_c$, and find that

$$1 - F(p) = \frac{2\sigma}{\sigma - 1}(p - p_c) + \mathcal{O}((p-p_c)^2) \tag{2.14}$$

as $p \to p_c$ from above. This implies that $P_\infty(p) \propto (p - p_c)$, and hence the mean-field critical exponent $\beta_{\text{perc}} = 1$.

2.2.3 MST connectedness functions

As noted above, the only difference between dynamic bond percolation and the Kruskal process is the necessity of enforcing the lack of cycles in the latter, and the only cycles on the wired Cayley tree are those that go out to the boundary. In order to calculate connectedness functions on the full

MST, it is therefore necessary to keep track of the values of p_0 at which different vertices become connected to the boundary. This leads us to define a "one-point function" for the MST problem as the probability $P_{(0;m)}(p_0)$ that the path on the MST from the origin 0 to the boundary has formed by time p_0 (*i.e.*, when all edges of cost $\leq p_0$ have been tested) and passes through a given vertex at distance m from the origin. In what follows, I use "path" to mean a self-avoiding walk that does not backtrack on itself.

Note that I set up this calculation on the wired Cayley tree and then assume no difficulties are introduced by taking the Bethe limit. As it turns out, the only M-dependent quantity involved is $F_M(p)$ from the percolation problem, which possesses a well-behaved limit.

We may relabel the vertices so that m is the origin and 0 is the specified vertex that the path to the boundary must pass through. For ever intermediate vertex $j = 1, \ldots, m-1$ along the path from the origin to the chosen vertex, there is a set of paths connecting j to the boundary through any of the $\sigma - 1$ side branches leaving this path. Define $\ell_{j,\infty}$ as the minimum over this set of paths of the maximum edge cost encountered along each of these paths (see figure 2.2): in other words, $\ell_{j,\infty}$ is the value of p at which a connection to the boundary is first formed along any of these $\sigma - 1$ side branches from vertex j in the dynamic bond percolation process. For $j = 0$ ($j = m$), define $\ell_{0,\infty}$ ($\ell_{m,\infty}$) as the minimum cost at which connection to the boundary is formed along any of the σ outward branches (resp., σ branches other than the specified path) in the Kruskal process.

Recursion for the one-point function

I then make use of the strong cycle property to come up with a recursion for $P_{(0;m)}(p_0)$. In going from $P_{(0;j-1)}$ to $P_{(0;j)}$ we add a set of cycles which go from 0 to j and through the side branches at j. In order to ensure that the edge connecting $j-1$ to j is on the MST, the most expensive edge on each of these cycles must fall somewhere on the side branch. If $X_{j-1}(p)$ is some set of conditions such that $P_{(0;j-1)}(p) = \Pr[X_{j-1}(p)]$, the recursion step is

$$P_{(0;j)}(p) = \Pr\left[\Big(\bigwedge_{i=0}^{j}(\ell_{i-1,i} < \ell_{j,\infty})\Big) \wedge (\ell_{j-1,j} < p) \wedge X_{j-1}(p)\right]. \tag{2.15}$$

where $\wedge$ denotes a logical 'and' and I simplify the notation by defining $\ell_{-1,0} = \ell_{0,\infty}$. Note that in determining whether the edge connecting $j-1$ to j is on the MST, we only need to impose $\ell_{i-1,i} < \ell_{j,\infty}$ for $j \geq i$. This fact allows the recursion to be readily solved: with the initial condition that we impose no direction for the path, $X_0(p) = (\ell_{0,\infty} < p)$ and the solution of the recursion is

$$P_{(0,m)}(p) = \Pr\left[\Big(\bigwedge_{j=1}^{m}\bigwedge_{i=0}^{j}(\ell_{j,\infty} > \ell_{i-1,i})\Big) \wedge \bigwedge_{j=0}^{m}(\ell_{j-1,j} < p)\right]. \tag{2.16}$$

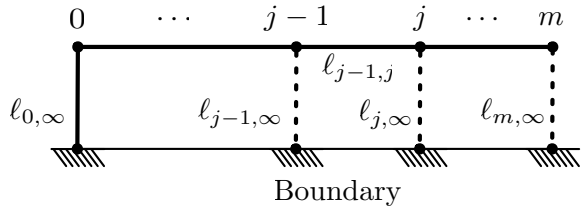

Figure 2.2: Construction of a path on the Bethe lattice starting from the origin (labeled m), passing through the vertex labeled 0, and connecting to the boundary, here shown as a shaded line. For this path to lie on the MST, the edge costs must obey the set of inequalities given in eq. (2.16). Note that each vertical line (both solid and dashed) stands for a subtree, not just a single edge, so this diagram depicts the entire Bethe lattice.

We can make (2.15) easier to work with by letting the origin have the same number of side branches as the intermediate vertices $j = 1, \ldots m-1$. I let $P'_{(0;m)}(p)$ denote this modification. Then (2.15) takes on a simple form if we work with the derivatives

$$\Phi_j(p) = \frac{d}{dp}P'_{(0;j)}(p), \tag{2.17}$$

representing the probability density that the connection described above is formed at p. The initial condition for the recurrence is

$$\Phi_0(p) = \frac{d}{dp}\big(1 - F(p)^\sigma\big). \tag{2.18}$$

All costs appearing in (2.15) are statistically independent, since they deal with disjoint sets of edges. We have

$$\Pr[\ell_{j,\infty} < p] = 1 - F^{\sigma-1}(p); \qquad \Pr[\ell_{j-1,j} < p] = p, \tag{2.19}$$

for $j = 1, \ldots, m$. The recurrence then becomes

$$\begin{aligned} \Phi_j(p) &= F(p)^{\sigma-1} \left(p\Phi_{j-1}(p) + \int_0^p dp'\, \Phi_{j-1}(p') \right) \\ &= F(p)^{\sigma-1} \frac{d}{dp} \left(p \int_0^p dp'\, \Phi_{j-1}(p') \right). \end{aligned} \tag{2.20}$$

The factor of $F^{\sigma-1}(p)$ is the probability that the $\sigma-1$ side branches from vertex j are not connected to the boundary by p. The two terms in the first line of (2.20) correspond to the two cases that the most expensive edge connecting the vertex j to the boundary on the path (which has cost p) respectively either is not or is the edge connecting vertices $j-1$ and j. In the former case, the edge connecting $j-1$ and j is already occupied and the connection takes place somewhere before it in the path with probability density $\Phi_{j-1}(p)$. In the latter case the rest of the path is already formed, with probability $P'_{(0,j-1)}(p)$, and the edge $j-1, j$ becomes occupied at p.

The recurrence may be written in terms of a kernel K:

$$\Phi_j(p) = \int_0^1 dp'\; K(p,p')\Phi_{j-1}(p'); \tag{2.21}$$

$$\begin{aligned} K(p,p') &= F(p)^{\sigma-1}\big[\theta(p-p') + p\delta(p-p')\big] \\ &= F(p)^{\sigma-1}\frac{d}{dp}\big[p\theta(p-p')\big], \end{aligned} \tag{2.22}$$

where $\theta(x)$ is the usual step function. This kernel has the meaning of a conditional probability density: it is the probability density (in the p variable) that j is first connected to the boundary along the specified path passing through $j-1$ at value p, given that $j-1$ is first connected to the boundary in the specified manner at value p'. Clearly, this must vanish if $p < p'$. If the factors of $F^{\sigma-1}$ are omitted, then we obtain the corresponding conditional probability for dynamic bond percolation, in which any number of connections to the boundary along side branches are allowed. Similar interpretations apply to the iterates of K that I consider below.

A related result phrased in terms of the Tutte polynomial was found by Steele and Fill in [87], although they did not mention the geometrical connection to percolation clusters.

Multi-point connectedness functions

All multi-point correlations on the wired Cayley tree may be built from the recurrence relation (2.21). Iteration of the recurrence requires iterated integrals of K:

$$\begin{aligned} K^{*m}(p_1, p_{m+1}) &\equiv K * K \cdots * K(p_1, p_{m+1}) \\ &= \int_0^1 dp_2 \cdots \int_0^1 dp_m \, K(p_1, p_2) \cdots K(p_m, p_{m+1}). \end{aligned} \tag{2.23}$$

The explicit step and δ- functions in each factor of K guarantee that $p_1 \geq p_2 \geq \cdots \geq p_{m+1}$. Referring to the definitions, we then have

$$P_{(0;m)}(p_0) = \int_0^{p_0} dp \, dp' \, F(p) K^{*m}(p, p') \Phi_0(p'), \tag{2.24}$$

where the factor $F(p)$ in the final integration restores the correct number ($\sigma + 1$) of branches at the origin.

We can obtain multipoint correlation functions by introducing branching into the path via additional chains of Ks. For example, the two-point connectedness function is the probability that two points at distances $m_1 > 0$, $m_2 > 0$ from a vertex that I label 0 are both connected to the boundary through 0 by the time $p = p_0$, and is given by

$$P_{(0;m_1,m_2)}(p_0) = \int_0^{p_0} dp_1 \, dp_2 \, dp_1' \, dp_2' \, F(p_1) F(p_2) K^{*m_1}(p_1, p_1') K^{*m_2}(p_2, p_2') \Phi_{(0;0,0)}(p_1', p_2') \tag{2.25}$$

(see figure 2.3). The initial distribution is $\Phi_{(0;0,0)}(p, p') = \delta(p - p') d/dp \; (1 - F^{\sigma-1}(p))$, because the connection of 0 to the boundary must not use either of the two side branches leading to m_1 and m_2. I then define the two-point correlation function $C^{(2)}(m, p_0)$ as the probability that two vertices at separation m are connected to each other and to the boundary by p_0. For $m > 0$, it is given by

$$C^{(2)}(m; p_0) = \sum_{m'=1}^{m-1} P_{(0;m',m-m')}(p_0) + 2P_{(0;m)}(p_0). \tag{2.26}$$

The term $2P_{(0;m)}(p_0)$ covers the cases where the path to the boundary on the tree from one vertex

passes through the other. For $m = 0$, I define $C^{(2)}(0; p_0) = P_{(0;0)}(p_0) = P_\infty(p_0)$.

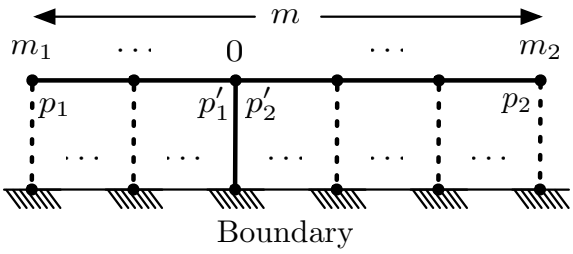

Figure 2.3: The set of paths contributing to $C^{(2)}(m; p_0)$, drawn using the same conventions as figure 2.2. Here I sum over all $m_1 \geq 0$, $m_2 \geq 0$ such that $m_1 + m_2 = m$.

This construction of the two-point connectedness function can be straightforwardly extended to multipoint correlations, giving the probability that, by time p_0, several specified vertices are on the same component of the tree and connected to the boundary. To compute the k-point connectedness function, one draws all possible trees connecting k leaf vertices to a root vertex 0, denoting the connection to the boundary. We associate a factor of F with each leaf vertex and an iterated kernel K^{*m_e} with each edge e of the tree, and sum over all the lengths $\{m_e\}$ subject to conditions on the separation of the leaf vertices. If n branches meet at the root vertex, we generalize $\Phi_{(0;0,0)}$ and associate with it a factor of

$$\Phi_{(0;0,\ldots,0)}(p_1, \ldots, p_n) = \frac{d}{dp_1} \left(1 - F(p_1)^{\sigma-1}\right) \prod_{j=2}^{n} \delta(p_1 - p_j). \tag{2.27}$$

Similarly, by comparing $d/dp_0 P_\infty(p_0)$ with $\Phi_0(p')$ and $\Phi_{(0;0,0)}(p'_1, p'_2)$, we may obtain the functions that must be associated with the other vertices of the tree: when one path splits into n paths we must insert a factor of

$$v_n(p_1; p_2 \cdots p_{n+1}) = \frac{\sigma + 1 - n}{\sigma + 1} \frac{1}{F(p_1)^n} \prod_{j=2}^{n} \delta(p_1 - p_j). \tag{2.28}$$

These k-point functions will be used in section 2.2.6 to analyze the geometry of the trees.

2.2.4 Asymptotic behavior of the iterated kernel

Eigenfunctions

To make further progress, we need to analyze the behavior of the iterated kernel K^{*m}. The analysis of $P_{(0,m)}(p_0)$ turns out to be simple: substituting (2.18) into (2.20) and using (2.12) shows that $\Phi_1(p) = \Phi_0(p)/\sigma$. In other words, Φ_0 is a right eigenfunction of K with eigenvalue $1/\sigma$. It follows that $\Phi_j(p) = \Phi_0(p)/\sigma^j$ for $j = 1, \ldots, m-1$, and

$$P_{(0,m)}(p_0) = \frac{1 - F(p_0)^{\sigma+1}}{(\sigma+1)\sigma^{m-1}} = \frac{P_\infty(p_0)}{(\sigma+1)\sigma^{m-1}}. \tag{2.29}$$

In particular, if we run the Kruskal process to completion by setting $p_0 = 1$, we have $P_{(0,m)}(1) = 1/(\sigma+1)\sigma^{m-1}$. The meaning of this result should be clear: due to the isotropy of the Bethe lattice, the path from the origin to infinity (which exists at p_0 with probability $P_\infty(p_0)$) must pass through one of the $(\sigma+1)\sigma^{m-1}$ vertices a distance m away, and all paths are equally probable, so the path is a random walk without backtracking on the Bethe lattice. In section 2.2.7, I will interpret this path as an isotropic random walk with fractal dimension $D_{\mathrm{p}} = 2$ in Euclidean space.

We can obtain all the right eigenfunctions of K by rewriting the eigenvalue equation for K, $\int_0^1 dp_2\, K(p_1, p_2) v_\lambda(p_2) = \lambda v_\lambda(p_1)$, in terms of the integral $V_\lambda(p) = \int_0^p dp'\, v_\lambda(p')$. This gives an ordinary differential equation for $V_\lambda(p)$,

$$\left(\lambda - pF(p)^{\sigma-1}\right) \frac{d}{dp} V_\lambda(p) = F(p)^{\sigma-1} V_\lambda(p), \tag{2.30}$$

whose general solution is

$$V_\lambda(p) = C \exp \int^p dp' \, \frac{F(p')^{\sigma-1}}{\lambda - p'F(p')^{\sigma-1}}. \tag{2.31}$$

for some proportionality constant C. By definition, V_λ must have the initial condition $V_\lambda(0) = 0$ and be continuous for all p. The only way to satisfy the inital condition (other than in the degenerate case $V_\lambda = 0$) is by having $\lambda = pF(p)^{\sigma-1}$ for some p, which implies that the eigenvalue λ lies in the interval $(0, 1/\sigma]$. For $0 < \lambda < 1/\sigma$, there are two solutions $p_\lambda^\pm$ to $\lambda = pF(p)^{\sigma-1}$, with

$p_\lambda^- = \lambda < p_c < p_\lambda^+$. The only way to satisfy continuity is by imposing $V_\lambda(p) = 0$ for all $p < p_\lambda^+$ (otherwise V_λ would be singular at p_λ^-). It may be verified that, for $\lambda = 1/\sigma$, the above solution reproduces the result we found previously: $p_\lambda^\pm = p_c = 1/\sigma$ and $v_{1/\sigma}(p) \propto \Phi_0(p)$.

As $p \to p_\lambda$ from above, the solutions (2.31) have a power law behavior,

$$V_\lambda(p) \sim (p - p_\lambda)^{\alpha_\lambda + 1}. \tag{2.32}$$

Expanding $F(p)$ near p_c, we find $\alpha_\lambda \sim -2\sigma(p_\lambda - p_c)$ as $p_\lambda \to p_c$ from above. As $p_\lambda^+ \to 1$ (*i.e.*, $\lambda \to 0$), we get $\alpha_\lambda \to -1$, and α_λ always lies between 0 and -1. This establishes that the eigenfunctions v_λ are positive and integrable for all $\lambda \in (0, 1/\sigma]$.

Green's function formalism

To make progress with multi-point connectedness functions such as (2.25), I first introduce the generating function

$$\widetilde{\Phi}_w(p) = \sum_{j=0}^{\infty} w^j \Phi_j(p), \tag{2.33}$$

for complex w. This may be thought of as a discrete Fourier-Laplace transform. Inserting this into (2.21) and summing over j yields

$$\int_0^1 dp' \, \Big[\delta(p - p') - wK(p, p')\Big] \widetilde{\Phi}_w(p') = \Phi_0(p). \tag{2.34}$$

I solve this integral equation by finding the resolvent operator (or Green function) g_w such that

$$\int_0^1 dp_2 \, \Big[\delta(p_1 - p_2) - wK(p_1, p_2)\Big] g_w(p_2, p_3) = \delta(p_1 - p_3). \tag{2.35}$$

The formal solution of (2.35) is simply

$$g_w(p_1, p_2) = \delta(p_1 - p_2) + \sum_{j=1}^{\infty} w^j K^{*j}(p_1, p_2). \tag{2.36}$$

Because the set of eigenvalues of K is the interval $(0, 1/\sigma]$, the series will not converge for $|w| > \sigma$, as we will see below. From the structure of K, we expect in general that g_w vanishes for $p_1 < p_2$. Because of the δ-function terms in (2.35) (note that K contains a δ-function piece), g_w contains a δ-function term as well as a smooth term. For $|w| < \sigma$, the δ-function part is clearly

$$g_w(p_1, p_2) = \frac{\delta(p_1 - p_2)}{1 - wp_1F(p_1)^{\sigma-1}} + \ldots, \tag{2.37}$$

where the omitted parts are ordinary functions. For w real and in the interval $[\sigma, \infty)$, the coefficient of the δ-function blows up at $p_1 = p_2 = 1/w$ and at $p_1 = p_2 = p_{1/w}$, in terms of the values p_λ^+ related to the eigenvalues λ of K as above.

Equation (2.35) can be converted into a differential equation by defining

$$G_w(p_2, p_3) = \int_0^{p_2} dp_2'\, g_w(p_2', p_3) \tag{2.38}$$

with the obvious boundary condition $G_w(0, p_3) = 0$. Then

$$\big(1 - wp_1F(p_1)^{\sigma-1}\big)\frac{d}{dp_1}G_w(p_1, p_2) - wF(p_1)^{\sigma-1}G_w(p_1, p_2) = \delta(p_1 - p_2). \tag{2.39}$$

For $p_1 \le p_c$ we have $F(p_1) = 1$ and (2.39) is solved by

$$G_w(p_1, p_2) = \frac{\theta(p_1 - p_2)}{1 - wp_1}, \tag{2.40}$$

which vanishes if $p_2 > p_c$.

The solution to (2.39) for all p_1 and p_2 is obtained by solving the equation without the δ-function term for $p_1 > p_2$ as in (2.30) (with $\lambda = 1/w$), and choosing the constant to obtain the correct discontinuous behavior at $p_1 = p_2$ (because of the δ-function). Because of the boundary condition $G_w(0, p_2) = 0$, we impose the requirement that G_w must vanish for $p_1 < p_2$ for all p_2. The result is

$$G_w(p_1, p_2) = \frac{\theta(p_1 - p_2)}{1 - wp_2F(p_2)^{\sigma-1}} \exp \int_{p_2}^{p_1} dp'\, \frac{wF(p')^{\sigma-1}}{1 - wp'F(p')^{\sigma-1}}. \tag{2.41}$$

Then g_w is obtained as $g_w(p_1, p_2) = \partial G_w(p_1, p_2)/\partial p_1$. Both g_w and G_w can be seen to be complex analytic for $w \notin [\sigma, \infty)$. g_w exhibits singular behavior if w is real and in $[\sigma, \infty)$ (as expected), and nowhere else, which means that the spectrum of K is precisely the interval $(0, 1/\sigma]$. For $w \notin [\sigma, \infty)$, there is no solution of the homogeneous equation obeying the initial condition that could be added to the solution. These results agree with the earlier statements for the two regimes $|w| < \sigma$ and $p_1 < p_c$.

Asymptotics of the iterated kernel

In order to obtain universal critical exponents, we need to investigate the scaling regime, which means probing the behavior of connectedness functions on length scales that are large relative to the lattice spacing, hence we need to determine the large m behavior of K^{*m}. Using the definition of g_w as a generating function, we have

$$K^{*m}(p_1, p_2) = \frac{1}{2\pi i} \oint dw \; \frac{g_w(p_1, p_2)}{w^{m+1}}, \tag{2.42}$$

where the contour is a small circle around the origin. Because g_w is analytic in w except on $[\sigma, \infty)$, we can increase the radius of the contour until we hit the first singularity in w, which is located on the real axis at the value w_c defined by $1 - w_c p_2 F(p_2)^{\sigma-1} = 0$. Parameterizing the contour as $w = w_c e^{-i\tau}$, the most strongly divergent piece of the integrand at this singularity will determine the large-m behavior of K^{*m} via

$$\int_{-\pi}^{\pi} d\tau \; \frac{e^{i\tau m}}{(i\tau)^z} = \frac{2\pi m^{z-1}}{\Gamma(z)} + i m^{z-1} \big(e^{-i\pi z} \Gamma(1 - z; -im\pi) - e^{i\pi z} \Gamma(1 - z; im\pi) \big), \tag{2.43}$$

where $\Gamma(z; t_0) = \int_{t_0}^{\infty} dt \; e^{-t} t^{z-1}$ is the incomplete gamma function. The first term on the right-hand side of (2.43) comes from the part of the contour coming in from $\tau = 0^+ + i\infty$, encircling the singularity at the origin and returning to $\tau = 0^- + i\infty$. If z is not an integer, there will be a branch cut from $\tau = 0$ to $\tau = +i\infty$, but we may use the same contour and obtain the same answer. The last terms come from the segments of the contour running from $\tau = -\infty$ to $\tau = -\pi$ and $\tau = \pi$ to

$\tau = \infty$; using the asymptotic form of the incomplete gamma function we see that

$$im^{z-1}\big(e^{-i\pi z}\Gamma(1-z;-im\pi) - e^{i\pi z}\Gamma(1-z;im\pi)\big) \sim \frac{2}{m\pi^z}\sin\left(\frac{2\pi m + \pi(1-z)}{2}\right) + \mathcal{O}(1/m^2), \quad (2.44)$$

so these terms may be neglected relative to the first as $m \to \infty$.

Exact asymptotic form of the iterated kernel

In order to obtain the explicit asymptotic behavior of K^{*m} in closed form, we need $F(p)$ in closed form, and therefore specialize to the case $\sigma = 2$ for the remainder of this section. In this case $F(p)$ was found in (2.13) and

$$\exp\int_{p_2}^{p_1} dp' \, \frac{wF(p')^{\sigma-1}}{1 - wp'F(p')^{\sigma-1}} = \left(\frac{p_2^w(1-w(1-p_1))}{p_1^w(1-w(1-p_2))}\right)^{\frac{1}{w-1}} \quad (2.45)$$

when $p_1 > p_2 > p_c$. Defining $h_w(p) = 1 - w(1-p)$ for brevity, we then have

$$g_w(p_1,p_2) = \frac{d}{dp_1}\left[\frac{\theta(p_1-p_2)}{h_w(p_2)}\left(\frac{p_2^w h_w(p_1)}{p_1^w h_w(p_2)}\right)^{\frac{1}{w-1}}\right]. \quad (2.46)$$

It might appear that (2.46) has an essential singularity at $w = 1$, but in fact

$$\lim_{w\to 1} g_w(p_1,p_2) = \frac{d}{dp_1}\left[\frac{\theta(p_1-p_2)}{p_1}e^{\frac{1}{p_1}-\frac{1}{p_2}}\right], \quad (2.47)$$

and the singularity with smallest $|w|$ is given by $h_{w_c}(p_2) = 0$; *i.e.* $w_c = 1/(1-p_2)$. With this definition g_w may be rewritten as

$$g_w(p_1,p_2) = \frac{d}{dp_1}\frac{\theta(p_1-p_2)}{(w_c-w)^{\frac{w}{w-1}}}\left(\frac{p_2^w h_w(p_1)}{p_1^w(1-p_2)^w}\right)^{\frac{1}{w-1}}. \quad (2.48)$$

Because $1/2 \leq p_2 \leq 1$, we know $1 \leq w_c/(w_c-1) \leq 2$ and the factor of $(w_c - w)$ is responsible for all divergences as $w \to w_c$. To find the leading-order divergence of this factor, we use

$$\begin{aligned}(w_c - w)^{-\frac{w}{w-1}} &= \exp\left[\log(w_c - w)\left(\frac{-w_c}{w_c - 1} - \sum_{n=1}^{\infty} \frac{(w_c - w)^n}{(w_c - 1)^{n+1}}\right)\right] \\ &= (w_c - w)^{-\frac{w_c}{w_c - 1}} + \mathcal{O}\left((w_c - w)^{-\frac{w_c}{w_c-1}+1} \log(w_c - w)\right), \end{aligned} \tag{2.49}$$

which, with $w = w_c e^{-i\tau}$, diverges as $(i\tau)^{-1/p_2}$. We may now use (2.43) to evaluate

$$K^{*m}(p_1, p_2) = \frac{1}{2\pi w_c^m} \int_{-\pi}^{\pi} d\tau \; e^{i\tau m} g_{w_c e^{-i\tau}}(p_1, p_2). \tag{2.50}$$

Because of cancellations that take place in (2.46) when $p_1 = p_2$, we must take the derivative with respect to p_1 explicitly before the $m \to \infty$ limit. The result is

$$\begin{aligned} K^{*m}(p_1, p_2) &\sim \delta(p_1 - p_2)(1 - p_2)^m \\ &\quad + \theta(p_1 - p_2) \frac{(1 - p_1)(1 - p_2)}{p_1(p_1 - p_2)^2} \left(\frac{p_2(p_1 - p_2)}{p_1(1 - p_2)}\right)^{1/p_2} \frac{(1 - p_2)^m \, m^{1/p_2 - 1}}{\Gamma(1/p_2)}. \end{aligned} \tag{2.51}$$

2.2.5 Asymptotics of the two-point function and mass

Building on the results of the previous section, I may now obtain the behavior of the two-point function $C^{(2)}$ and hence the scaling properties of the mass of one component of the Bethe lattice MSF. I present two versions: an essentially exact version for $\sigma = 2$, and an asymptotic calculation valid for all σ in the region $p_0 - p_c$ small.

Exact asymptotics

If we make no further approximations, when we calculate correlation functions such as (2.26) it is clearly easier to integrate over p_2 first and then find the large-m behavior rather than attempt to integrate (2.51) directly over p_2. As an example we now calculate the asymptotic behavior of

$C^{(2)}(m;p_0)$. Inserting δ_{m,m_1+m_2} in (2.26) gives

$$C^{(2)}(m;p_0) = \int_0^{p_0} dp_1 dp_2 \; \frac{1}{2\pi w_c^m} \int_{-\pi}^{\pi} d\tau \; e^{i\tau m} \widetilde{\Phi}_{w_c e^{-i\tau}, w_c e^{-i\tau}}(p_1,p_2), \tag{2.52}$$

where

$$\widetilde{\Phi}_{w_1,w_2}(p_1,p_2) = \int_0^1 dp_1' dp_2' \; g_{w_1}(p_1,p_1') g_{w_2}(p_2,p_2') \Phi_{(0;0,0)}(p_1',p_2'). \tag{2.53}$$

with $\Phi_{(0;0,0)}(p_1',p_2') = \theta(p_1'-p_c)\delta(p_1'-p_2')/{p_1'}^2$ for $\sigma = 2$. This integral is done easily:

$$\widetilde{\Phi}_{w,w}(p_1,p_2) = \\ \frac{d}{dp_1}\frac{d}{dp_2}\frac{1}{w+1}\left(\frac{h_w(p_1)h_w(p_2)}{p_1^w p_2^w}\right)^{\frac{1}{w-1}} \left[\left(\frac{1}{2-w}\right)^{\frac{w+1}{w-1}} - \left(\frac{p}{h_w(p)}\right)^{\frac{w+1}{w-1}}\right], \tag{2.54}$$

where $p = \min(p_1,p_2)$. Clearly, since the integrand in (2.53) has no support below p_c, we obtain a branch cut starting at $w_c = 1/p_c = 2$ and this is the maximum radius the contour in w may take. Since the only divergence as $w \to 2$ comes from the first term in the square brackets, we obtain

$$\widetilde{\Phi}_{2e^{-i\tau},2e^{-i\tau}}(p_1,p_2) = \frac{d}{dp_1}\frac{d}{dp_2}\frac{h_2(p_1)h_2(p_2)}{p_1^2 p_2^2}\left[\frac{1}{3}\left(\frac{1}{2i\tau}\right)^3 + \mathcal{O}\left(\frac{\log\tau}{\tau^2}\right)\right]. \tag{2.55}$$

Then, using

$$\int_{-\pi}^{\pi} d\tau \; \frac{e^{i\tau m}}{(i\tau)^3} \sim m^2\pi + \mathcal{O}(1/m) \tag{2.56}$$

and

$$\int_{p_c}^{p_0} dp \; \frac{d}{dp}\frac{h_2(p)}{p^2} = \frac{2p_0-1}{p_0^2} \tag{2.57}$$

we finally obtain

$$C^{(2)}(m;p_0) \sim \left(\frac{2p_0-1}{p_0^2}\right)^2 2^{-m}\left(\frac{m^2}{48} + \mathcal{O}(m\log m)\right). \tag{2.58}$$

We see that $C^{(2)}(m;p_0)$ scales as m^2/σ^m for large m, and this behavior is obtained for *any* $p_0 p_c$.

Asymptotics in the critical regime

Next I present an alternative calculation valid for all σ when $p_0 - p_c$ is small. I will define $\delta p = p - p_c$ (but *e.g.* $\delta(p_1 - p_2)$ is a δ-function as usual). First we write

$$pF(p)^{\sigma-1} \simeq p_c - \delta p, \tag{2.59}$$

$$F(p)^{\sigma-1} \simeq 1 - 2\sigma\delta p, \tag{2.60}$$

valid as $p \to p_c$ from above. Then, using (2.41), G_w can be calculated for $w \notin [\sigma, \infty)$. Because all the singular behavior as $w \to \sigma$ arises from the denominator of the integrand of (2.41), we may approximate $F^{\sigma-1}(p) \simeq 1$ in the numerator and

$$G_w(p_1, p_2) = \theta(p_1 - p_2)\frac{\delta p_1 + 1/w - p_c}{w(\delta p_2 + 1/w - p_c)^2} \tag{2.61}$$

for δp_1, δp_2 both small and positive. Then if we consider the transform

$$\widetilde{C}^{(2)}_w(p_0) = \sum_{m=0}^{\infty} w^m C^{(2)}(m, p_0) \tag{2.62}$$

we notice that the transform of the sum over m_1 in eq. (2.26) becomes simply a product of g_ws inside the integral in eq. (2.25), and we neglect the factors $F(p_1)$, $F(p_2)$ as these do not affect the leading m dependence. Further we can neglect $2P_{(0;m)}$ as it falls off faster than the term we keep. Also $\Phi_{(0;0,0)}(p_1, p_2) = 2\sigma\delta(p_1 - p_2)$ for p_1 above p_c, and zero below. Finally, we will estimate the expected mass inside radius m,

$$\overline{M(m, p_0)} = C^{(2)}(0; p_0) + \sum_{m'=1}^{m} (\sigma + 1)\sigma^{m'-1} C^{(2)}(m', p_0) \tag{2.63}$$

directly, as this is similar to the definition of the transform of $C^{(2)}$: we simply evaluate the transform at $w = \sigma(1 - 1/m)$, which cuts off the sum at around m. This is a value at which G_w is not

singular. Thus we have to calculate

$$\widetilde{C}_w^{(2)}(p_0) = 2\sigma(\delta p_0 + 1/w - p_c)^2 \int_0^{\delta p_0} d\delta p' \, \frac{1}{w^2(\delta p' + 1/w - p_c)^4}. \tag{2.64}$$

For fixed p_0, the dominant contribution comes from the lower limit, and contains $(1-wp_c)^{-3}$ times factors that go to constants as $w \to \sigma = 1/p_c$. Hence the mass behaves as

$$\overline{M_{\rm soft}(m, p_0)} \sim \frac{2}{3}\sigma(\sigma+1)\delta p_0^2 m^3 \tag{2.65}$$

as $m \to \infty$ [we inserted a factor $(\sigma+1)/\sigma$ to account for the number of neighbors $\sigma+1$ at the first step in $\overline{M(m,p_0)}$, as in eq. (2.63)]. This method in fact differs from the definition above in using a soft cut-off for the sum over m' instead of a hard one, $m' \leq m$. Now that the form of the summand is known, we can evaluate it using either form of cutoff. Hence we find that for the *hard* cut-off, the result is smaller by a factor 6:

$$\overline{M(m, p_0)} \sim \frac{1}{9}\sigma(\sigma+1)\delta p_0^2 m^3. \tag{2.66}$$

Thus the correlation function behaves as

$$C^{(2)}(m, p_0) \sim \frac{1}{3}(\sigma\delta p_0)^2 m^2 \sigma^{-m}, \tag{2.67}$$

which agrees with the result for $\sigma = 2$.

Remarks

Equations (2.66) and (2.67) are the main results of this section. From the structure of the expression for G_w, the lower limit always dominates, so this m dependence holds for all $p_0 > p_c$, at sufficiently large m. More precisely, the results are valid only if δp_0 is greater than order $1/m$. This means that m is much larger than the correlation length at p_0, which is proportional to $1/(\delta p_0)$. The main contribution to the integral is from $\delta p'$ less than of order $1/m$. This is in agreement with the "superhighways" idea [20, 22], which I will return to in section 3.6.1.

Notice also that the factor δp_0^2 is present in both results because the probability that one vertex is connected to infinity is $P_\infty(p_0) \simeq 2\sigma(\sigma+1)\delta p_0/(\sigma-1)$, for δp_0 small and positive, and so is $\propto \delta p_0^2$ for two vertices (the two events are uncorrelated, because the two points are separated by more than the correlation length). The dependence on σ is the same within subleading terms of relative order $1/\sigma$.

2.2.6 k-point connectedness functions and moments of the mass

Definitions

We may extend the method to estimate asymptotics of correlation functions of any order k, that is the probability $C(i_1, \ldots, i_k; p_0)$ that some given set of vertices (labeled i_1, i_2, ..., i_k) are on the same connected component of MSF(p_0) and that this component is infinite. The procedure is straightforward: to compute the k-point correlation function for a given set of k distinct vertices, one draws the smallest subtree of the Bethe lattice such that all k given vertices are connected, and choose any vertex on this subtree as the root point, along which the connection to infinity occurs in the MSF(p_0) (eventually, we will sum over the possible root points). Thus the leaves of the subtree (*i.e.* the vertices of degree one) must be among the given k vertices, but if any of the k given vertices are not leaves, they can be anywhere on the subtree. Starting from the root, we propagate out to (or possibly through) each of the k given vertices, along the subtree. The subtree can be viewed as made of chains of edges connected by degree-two vertices, with the ends of the chains at either (i) the root point, which has degree ≥ 1, (ii) the leaves of the subtree, or (iii) vertices of degree > 2 other than the root point. For each such chain e of m_e steps, we associate the iterated kernel $K^{*m_e}(p_i, p_j)$, where the labels i, j are associated to the two ends of the chain, with i the end further from the root point. For the initial distribution at the root, if there are n chains leaving it ($n \leq \sigma$), we generalize $\Phi_{(0;0,0)}$ to

$$\Phi_{(0;0,\ldots,0)}(p_1, \ldots, p_n) = \left[\frac{d}{dp_1}\left(1 - F^{\sigma+1-n}(p_1)\right)\right] \prod_{j=2}^{n} \delta(p_1 - p_j). \tag{2.68}$$

Similarly, by comparing $dP_\infty(p_0)/dp_0$ with $\Phi_0(p')$ and $\Phi_{(0;0,0)}(p'_1, p'_2)$, we see that at a vertex of the subtree of degree $n \neq 2$, we must associate with it a factor

$$v_n(p_1, p_2 \ldots, p_n) = F^{2-n}(p_1) \prod_{j=2}^{n} \delta(p_1 - p_j). \tag{2.69}$$

After multiplying together all these factors, we must integrate over all the parameters like p_i between the limits 0 and p_0. There are two of these parameters for each chain on the subtree; clearly some could be eliminated using the δ functions. Finally, we must sum over all possible root points on the tree. This procedure yields $C(i_1, \ldots, i_k; p_0)$ (unlike the earlier description for the $k = 2$ case, there are no exceptions to this prescription for cases of vertices coinciding with each other or with the root point).

Higher moments of the fractal mass

The higher-point correlation functions can be used to calculate higher moments of the mass, $\overline{M(m, p_0)^k}$. These are the average of the kth power of the sum over positions at distance less than m from the origin of the "indicator function" that is one if and only if the vertex is on the same connected component of MSF(p_0) as the origin. $\overline{M(m, p_0)^k}$ is equal to the sum of the $k+1$-point correlation function $C(i_1, \ldots, i_{k+1}; p_0)$ over all positions of i_2, ..., i_{k+1} within m steps of the origin at i_1.

For m large, the largest contribution to $\overline{M(m, p_0)^k}$ will come from configurations of i_l, ($l = 1$, ..., $k+1$) for which, in the subtree in the calculation of $C(i_1, \ldots, i_{k+1}; p_0)$, all the given vertices are at its leaves, the root point has degree 2, and the vertices of degree > 2 have degree 3. For these there are $2k$ chains (iterated kernels) in the subtree. We can estimate the power of m as $m \to \infty$ using the same approximations as for $k = 1$. The sums over position are estimated by using the propagator g_w in place of all K^{*m_e}'s, with $w = \sigma(1 - 1/m)$ in each one. The factors at the vertices of the subtree (other than δ-functions) can be dropped, at least when δp_0 is small (and for larger δp_0 do not affect the scaling behavior). The integrals over the p_i associated with the leaves can be done by using G_w in place of g_w for these chains. The remaining integrals over p_i's associated with the other vertices and the root point are dominated by the lower limit $\delta p_i = 0$, and can be estimated

by power counting. As each additional leaf on the subtree leads to an extra factor $G_w g_w$ and one additional integral similar to that at the root (as for $k = 1$ above), this yields finally (neglecting constant factors)

$$\overline{M(m,p_0)^k} \sim m^{3k} \sim [\overline{M(m,p_0)}]^k. \tag{2.70}$$

As all kth moments scale like the kth power of the first moment, this means that the (random) mass $M(m,p_0)$ does not have a very broad distribution, and its typical behavior is well-described by its expected value. Hence the connected components of MSF(p_0) on the Bethe lattice are not multifractals.

2.2.7 Fractal dimensions and mean-field theory

So far I have developed a method for computing correlation functions on the Bethe lattice, while our real interest is in lattices in Euclidean space of dimension d. For sufficiently high d, we would expect that a mean field theory holds for quantities such as exponents; this assertion will be justified *post hoc* in the following chapter by a perturbation analysis of corrections due to fluctuations neglected in the mean field theory. The Bethe lattice results provide the mean-field theory results, once we have explained how to convert them to apply to Euclidean space.

For a hypercubic lattice on Euclidean space, if we choose a path starting from the origin randomly (with equal probability for each), then it behaves as a random walk, and after m steps will be of order $\sqrt{m}$ in Euclidean distance from the origin. In the set of all paths from the origin, any two paths initially coincide but ultimately part company. If we neglect the possibility that they subsequently intersect (and also that a path may intersect itself, including by backtracking), then the union of the paths forms a tree, equivalent to the Bethe lattice with $z = 2d$. Hence in this correspondence, separations on the Bethe lattice behave like distances *squared* on the Euclidean lattice [85],

$$m \sim r^2 \tag{2.71}$$

which allows us to infer mean-field scaling dimensions from the Bethe lattice theory. More formally, to apply the Bethe lattice results as an approximation for the lattice, in equation (2.20) we

must now sum over all neighbors of the given site, since the path may go through the site in any direction. Equation (2.33) needs to be replaced by a Fourier transform

$$\widetilde{\Phi}_{\mathbf{k}}(p) = \sum_{\mathbf{x}} e^{i\mathbf{k}\cdot\mathbf{x}} \Phi_{\mathbf{x}}(p), \tag{2.72}$$

which means that in (2.34) and all subsequent equations we make the substitution

$$w \to \sum_{\hat{n}_j} e^{-i\mathbf{k}\cdot\hat{n}_j} = 2d - \mathbf{k}^2 + O(|\mathbf{k}|^4), \tag{2.73}$$

where $\{\hat{n}_j\}$ are the basis vectors of the hypercubic lattice. This leads to the same relation (2.71).

We established that the probability the path from any vertex to infinity passes through a given vertex a distance m away behaves as $P_{(0;m)}(p_0) \sim \sigma^{-m}$; summing this over all sites within a ball of radius m on the Bethe lattice gives

$$P_{(0;0)}(p_0) + \sum_{m'=1}^{m} (\sigma+1)\sigma^{m'-1} P_{(0;m')}(p_0) \sim m. \tag{2.74}$$

Then we expect that in Euclidean space, the mass of the path lying within radius R scales as $M_p(R, p_0) \sim R^2$, consistent with the picture of this path (mentioned earlier for the Bethe lattice) as a random walk, with dimension $D_{\mathrm{p}} = 2$. This is the same as the dimension of the backbone of a critical percolation cluster for $d \geq d_c = 6$ [39].

Similarly, the mass of the component connected to the origin on the Bethe lattice, $M(m, p_0) \sim m^3$ becomes $M(R, p_0) \sim R^6$ within radius R in Euclidean space, meaning that a connected component of the MST has a mean-field fractal dimension $D = 6$. Because the union of the spanning trees fills the lattice, this strongly suggests that the critical dimension of the MST will be $d_c = 6$, as discussed in Section 2.1.6. This is the same critical dimension as for percolation (at threshold). We emphasize again that the result is valid for p_0 greater than of order $1/m \sim 1/R^2$. Since the correlation length ξ behaves as $|p_0 - p_c|^{-\nu_{\mathrm{perc}}}$, with $\nu_{\mathrm{perc}} = 1/2$ for $d > 6$, this means it holds for $R > \xi$, and thus involves distance scales at which ordinary correlations for percolation decay exponentially.

2.3 Related applications

In this section I briefly discuss the applicability of the scaling results obtained in the previous section to other models of interest — this material will not be needed to understand the arguments that follow. In section 2.3.1 I first consider a simple extension of the wired Bethe lattice model examined above, in which we allow the degree of each vertex to be an iid random variable. Since the geometry of these graphs is still locally treelike, with the only cycles present being infinite (going out to the boundary), it should not be surprising that my previous analysis carries through essentially unchanged.

In section 2.3.2 I address the continuum MST problem, briefly referred to in 1.1.1 — recall that in this case, the edge costs are distances induced by an embedding of the vertices in Euclidean space. The appropriate mean-field approximation to such a model is the Poisson-weighted infinite tree (PWIT). This is another example of a locally treelike graph, so my analysis is also applicable here.

2.3.1 Random, locally treelike graphs

The results of the previous sections admit a simple generalization to certain classes of random graphs. We consider a tree where the coordination number of each vertex is an iid random variable distributed according to $\rho(\sigma)$. In what follows, averages with respect to ρ are denoted by angle brackets. The fixed-point equation for $F(p)$, (2.12), becomes

$$F(p) = 1 - p + p \sum_{\sigma=0}^{\infty} \rho(\sigma) F(p)^{\sigma}. \tag{2.75}$$

Again, $F(p)$ is defined as the smallest solution at fixed p. In what follows we assume $1 < \langle\sigma\rangle < \infty$, which are the conditions necessary for the graph to admit a conventional percolation transition: $F(p) = 1$ for $p \leq p_c = 1/\langle\sigma\rangle$, $F(p) < 1$ for $p > p_c$, and we can construct a nontrivial ensemble of MSTs under wired boundary conditions at infinity.

At this point we find it helpful to multiply the kernel defined in (2.22) by a factor of σ. This has the effect of summing the connectedness functions over lattice sites (done in, *e.g.*, (2.63))

simultaneously with averaging over edge costs. This modification allows us to relate the spectrum and eigenfunctions of the random graph kernel $\langle \sigma K \rangle$ back to those obtained above: *e.g.*, $\langle \Phi_0(p) \rangle$ is an eigenfunction of $\langle \sigma K \rangle$ but *not* of $\langle K \rangle$.

The entire calculation goes through as before, with occurrences of $F(p)^{\sigma-1}$ in the Bethe lattice Green's function replaced by $\langle \sigma F(p)^{\sigma-1} \rangle$ and the transform variable w rescaled by $1/\langle \sigma \rangle$. Under the additional assumption $\langle \sigma^2 \rangle < \infty$, one may repeat the asymptotic analysis of section 2.2.5 and obtain

$$\left\langle \overline{M_{\text{soft}}(m, p_0)} \right\rangle \sim \frac{2}{3} \langle \sigma(\sigma+1) \rangle \delta p_0^2 m^3. \tag{2.76}$$

The generalization extends to the higher moments of the cluster mass discussed in section 2.2.6.

One may also consider "quenched" moments of cluster masses, of the form

$$\mathcal{M}_{k,\ell}(m, p_0) \equiv \left\langle \left(\overline{M(m, p_0)^k} \right)^\ell \right\rangle \tag{2.77}$$

for $\ell > 1$. These quantities cannot readily be calculated with the techniques discussed above and are beyond the scope of this section.

2.3.2 Poisson-weighted infinite tree

Another possible specialization of the MST problem (which is frequently relevant in practice) is the continuum model introduced in section 1.1.1, in which the edge costs are distances between the corresponding vertices, here taken to be distributed according to a Poisson process on $\mathbb{R}^d$. This means the underlying graph is formally the infinite complete complete graph.

An appropriate mean-field model for this variant of the MST problem is the Poisson-weighted infinite tree (PWIT) [15, 44, 45], which effectively ignores all triangle inequality-type correlations induced by the metric structure of $\mathbb{R}^d$ and keeps only the probability distribution for separations of points. The PWIT is a tree with infinite degree at each vertex, and the costs associated with the edges incident on each vertex are iid and given by a Poisson process on $\ell \geq 0$ with density $\rho(\ell) \propto \ell^{d-1}$, which is the same as the measure on the set of distances between the points of the uniform Poisson process on $\mathbb{R}^d$.

Since the PWIT is a tree, we may view it as the $\sigma \to \infty$ limit of the Bethe lattice used in the calculations above. As we have noted above, the form of the MST depends only on the relative order of edge costs and not on the distribution they are drawn from. The results on the Bethe lattice do not have a sensible $\sigma \to \infty$ limit, so it is necessary to produce a finite-degree random tree by deleting the highest-cost edges. In order to obtain an average degree $\sigma + 1$, we cut off the distribution with $\rho(\ell)$ at $\ell = \mathcal{O}((\sigma+1)^{1/d})$. For each $d \geq 1$ in this model, the percolation threshold ℓ_c is non-zero, and since the behavior is dominated by edges close to ℓ_c, the results for the PWIT will be in the same universality class as the Bethe lattice model.

In fact, if we identify $\ell = (\sigma + 1)p$, the $\sigma \to \infty$ limit of the Bethe lattice model is the PWIT for $d = 1$. Notice that the limit of our expressions exists, because by writing $F(p) = 1 - G(\ell)/\sigma$, then $F^{\sigma-1} \to e^{-G}$, and the fixed-point equation for F becomes [45]

$$G(\ell) = \ell(1 - e^{-G(\ell)}). \tag{2.78}$$

In the transforms, we should also set $s = w/\sigma$. Then the limit of the theory makes sense for the masses M, M_p within m steps of the origin: $\overline{M} \sim \frac{2}{3}\delta\ell^2 m^3$ (again, $\delta\ell = \ell - \ell_c$, where here $\ell_c = 1$). For the PWIT one would naturally wish to express such quantities in terms of the distance ℓ defined as the sum of the ℓ_e's along a path on the tree, and since most edges accepted are either below, or not far above, ℓ_c, these masses $M(\ell)$, $M_p(\ell)$ scale the same way, and again $\ell \sim R^2$ because the paths are random walks. When this measure of distance is used, the number of connected components inside ℓ is also well-behaved. Note that because it is believed that continuum percolation is in the same universality class as bond percolation, we would expect universal properties of the continuum MST to be the same as those of the lattice model anyway, so that all of these conclusions are consistent.

2.4 Conclusions

In this chapter, I have achieved the following results. For a finite graph, I defined a process MSF(p_0) which is a random forest that becomes the MST for $p_0 = 1$. Using the Bethe lattice with

wired boundary conditions, and taking the infinite size limit, I showed that the infinite connected components of MSF(p_0) on the Bethe lattice contain of order m^3 vertices within m steps of any vertex on this component, for any p_0 greater than the value p_c of the threshold for bond percolation on the Bethe lattice. This result is essentially rigorous. Transferring it (heuristically) to Euclidean space, this means that the mass of an infinite connected component of MSF(p_0) within a ball of radius R scales as R^D with $D = 6$, for d sufficiently large and p_0 greater than the value p_c of the threshold for bond percolation on the lattice used. This then implies that for $d > 6$ there are of order R^{d-6} large connected components that intersect such a ball. I also gave a non-rigorous second argument for these results, using scaling ideas (this argument directly addresses the critical dimension d_c above which the results hold). The results also hold (rigorously) for the Poisson-weighted infinite tree, and (heuristically) for the continuum MST model in Euclidean space.

Following the reasoning of Newman and Stein [23], these results for the MST imply that the strongly-disordered spin-glass model has an uncountable number of ground states for $d > 6$, of which of order $2^{\mathcal{O}(R^{d-6})}$ can be distinguished within a ball of radius R. For $d \leq 6$, the logarithm of the number of ground states is smaller than any power of R, and possibly only of order one, or simply one (with probability one).

Chapter 3

Minimum spanning trees: Loop corrections

In the previous chapter I developed a mean-field theory for connectedness functions of components of the minimum spanning tree (MST), by working on the Bethe lattice. The corresponding fractal dimensions were argued (in section 2.2.7) to be valid for the MST on Euclidean lattices of dimension greater than d_c. In this chapter I examine how these results are modified for $d < d_c$, for the specific case $p_0 \leq p_c$. Spanning clusters do not proliferate in low spatial dimension (recall section 2.1.2), so within a window we expect to see of order one space-filling MSF component. This means that the k-point connectedness functions defined earlier have trivial behavior ($D_{\text{MSF}} = d$). In this chapter, I therefore restrict my attention to the "one-point" function, or alternatively the fractal dimension D_p of paths on the MST.

On the Bethe lattice, the only difference between the Kruskal process and dynamic bond percolation occurred when a growing cluster formed a cycle which extended out to the (wired) boundary of the lattice. In this chapter, I describe how to account in a consistent manner for the fact that clusters on a finite dimensional lattice will also develop cycles of finite extent, which must also be dealt with in accordance with Kruskal's algorithm. For simplicity, in this chapter I neglect the presence of paths to the wired boundary, which means I investigate MSF(p_0) (defined in section 2.1.3) for $p_0 \leq p_c$: in other words, minimal spanning forests (MSFs) on subcritical percolation

clusters. The implications of this will be discussed in section 3.6.

I accomplish this by, again, making use of the close correspondence between the Kruskal process and dynamic bond percolation (section 2.1.3). Therefore, I begin in section 3.1 by reviewing low-density expansions for percolation connectedness functions on finite-dimensional graphs of finite extent. I formulate this expansion in terms of basic geometric notions such as the existence of paths. In section 3.2 I implement the no-cycle condition to obtain a lattice expansion for connectedness functions on components of MSF(p_0). Although the primary utility of these results is to provide a basis for the continuum field theory of section 3.3, the lattice expansion is of interest in its own right and may be taken as the starting point for a mathematically rigorous formulation of these results using *e.g.*, lace expansion methods [49].

In section 3.3 I describe how to take the continuum limit of this lattice expansion, obtaining a perturbative expansion in terms of Feynman diagrams. I emphasize that my result differs from conventional field theory in that my results are formulated entirely in terms of a perturbative, diagrammatic expansion; at no point do I make reference to an action as a generating functional for these diagrams. Because the MST problem is nonlocal, I find it likely that no such action defined in terms of a finite number of locally-interacting fields exists (I will elaborate on this point in section 3.7).

In section 3.4 I prove that a consistent renormalization group (RG) may be defined for my perturbation expansion. This is a highly non-trivial task (and the section in question is quite technical), for the reasons just explained: my expansion is defined in terms of performing various manipulations on percolation diagrams, and I must prove various factorization properties to establish that these manipulations are sufficiently local for the RG to be defined in any sensible way. I find that the theory is renormalizable to all orders — this is perhaps the most remarkable result of this chapter. In section 3.5 I then perform a one-loop RG calculation and obtain the fractal dimension of paths on MSF(p_0) as an asymptotic series in $\varepsilon = d_c - d$, which is the main quantitative result of the chapter.

3.1 Graphical expansions for percolation

In this section I review the construction of low-density (high-temperature) graphical expansions for connectedness functions in bond percolation. Bond percolation on a finite-dimensional lattice is conventionally treated as the $Q \to 1$ limit of the low-density expansion of the Q-state Potts model [88–90]. The partition function this model is a polynomial in Q and also arises as the generating function for Q-colorings of the vertices of the graph upon which the model is formulated, with weights that depend on whether adjacent vertices are given the same or different colors. This generating function is the celebrated Tutte polynomial of graph theory [91]. Although the $Q \to 1$ limit of this function can be taken *a posteriori*, it lacks a mathematical definition in terms of state variables (colors), so in order to establish a correspondence with percolation I instead use a method originally due to Essam [40, 92]. This has the advantage of being phrased explicitly in terms of geometric quantities, and is closely related to the lace expansion [49] from rigorous statistical mechanics.

I summarize the relevant results from percolation in this section both for the purposes of introducing notation and terminology, and also because my discussion of MSFs in the following section will closely parallel the computations done here. In section 3.1.1 I define an expansion (3.8) for two-point connectedness functions, with the generalization to n-point functions being given in section 3.1.2. In section 3.1.3 I show that these are in fact equivalent to the usual Potts model expressions. In section 3.1.4 I take the first step towards a continuum theory by introducing the notion of "topological graphs," defined independently of an embedding in the lattice. The main results of this section are the graphical expansions (3.8), (3.16) with diagrammatic weights given in simplest form in (3.13), (3.18).

3.1.1 Two-point connectedness functions

Essam's expansion for percolation [40, 92] is based on the principle of inclusion and exclusion from elementary probability theory [93]. This may not be familiar to readers with a physics background, so I summarize it here. We start with a set of events (*i.e.*, boolean-valued functions) $\{X_i\}$

for i in some index set I. In order to calculate probabilities we introduce the indicator function

$$\mathbb{I}[X_i] \equiv \begin{cases} 1 & X_i \text{ true,} \\ 0 & X_i \text{ false.} \end{cases} \tag{3.1}$$

The probability that an event X_i will happen is then simply the expectation value $\langle \mathbb{I}[X_i] \rangle$ taken with respect to the appropriate ensemble. The principle of inclusion-exclusion is the expansion

$$\mathbb{I}\Big[\bigvee_{i\in I} X_i\Big] = \sum_{\emptyset \neq I' \subseteq I} (-1)^{|I'|+1} \mathbb{I}\Big[\bigwedge_{j\in I'} X_j\Big]. \tag{3.2}$$

where '$\vee$' denotes a logical 'or' and '$\wedge$' denotes a logical 'and'. An analogous series may be obtained for the probability that all the events occur, by using De Morgan's law $\neg(\vee_i X_i) = \wedge_i(\neg X_i)$, where '$\neg$' denotes logical negation. This yields

$$\mathbb{I}\Big[\bigwedge_{i\in I} X_i\Big] = \sum_{I' \subseteq I} (-1)^{|I'|} \mathbb{I}\Big[\bigwedge_{j\in I'} \neg X_j\Big]. \tag{3.3}$$

Note that in this case I' may be the empty set, so the first term of this series is 1.

We apply this to bond percolation at a parameter value p by first investigating the two-point connectedness function, defined as

$$\begin{aligned} C_{\mathbf{x},\mathbf{y}}(p) &\equiv \langle \mathbb{I}[\mathbf{x}, \mathbf{y} \text{ connected at } p] \rangle \\ &= \langle \mathbb{I}_c(\mathbf{x}, \mathbf{y})|_p \rangle. \end{aligned} \tag{3.4}$$

where the angle brackets denote an average with respect to all realizations of the edge costs. The points $\mathbf{x}, \mathbf{y}$ are connected if at least one self-avoiding walk (referred to as a "path" in what follows) on the lattice connecting them has formed by the time the percolation parameter has been raised to ℓ_0. Defining $\Gamma_{\mathbf{x},\mathbf{y}}$ to be the set of all self-avoiding walks on the lattice between $\mathbf{x}$ and $\mathbf{y}$, we may write

$$\mathbb{I}_c(\mathbf{x}, \mathbf{y})|_p = \mathbb{I}\Big[\bigvee_{\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}} (\gamma \leq p)\Big]. \tag{3.5}$$

where I define the event

$$(\gamma \leq p) \equiv \left(\max_{e \in \gamma} \ell_e \leq p \right); \tag{3.6}$$

i.e., we require all edges e on the path γ to be present by the time the parameter is raised to the value p. Using equation (3.2) to expand the right-hand side of (3.5) by inclusion-exclusion yields

$$\mathbb{I}_c(\mathbf{x}, \mathbf{y})|_p = \sum_{\emptyset \neq \Gamma' \subseteq \Gamma_{\mathbf{x},\mathbf{y}}} (-1)^{|\Gamma'|+1} \mathbb{I}\left[\bigwedge_{\gamma' \in \Gamma'} (\gamma' \leq p) \right]. \tag{3.7}$$

I obtain an expansion in terms of graphs from (3.7) by grouping together all terms that test the same set of edges on the lattice; the terms in the series are now indexed by graphs whose edges are "covered" by subsets Γ' of the paths in the original summation, in the sense that every edge on the graph must be on at least one path. Because the paths are self-avoiding walks, all the graphs $\mathcal{G}_{\mathbf{x},\mathbf{y}}$ generated from their unions must be vertex-irreducible: removing any vertex from the graph must leave at least one of the points $\mathbf{x}, \mathbf{y}$ in each connected component. I denote the set of such graphs by $\mathcal{G}_{\mathbf{x},\mathbf{y}}$.

Letting $\Gamma_{\mathbf{x},\mathbf{y}}(G \in \mathcal{G}_{\mathbf{x},\mathbf{y}})$ denote the set of paths on G connecting the root points $\mathbf{x}, \mathbf{y}$, equation (3.7) becomes

$$\mathbb{I}_c(\mathbf{x}, \mathbf{y})|_p = \sum_{G \in \mathcal{G}_{\mathbf{x},\mathbf{y}}} \sum_{\Gamma' \subseteq \Gamma_{\mathbf{x},\mathbf{y}}(G)} (-1)^{|\Gamma'|+1} \mathbb{I}[\Gamma' \text{ covers } G] \mathbb{I}[G \leq p]. \tag{3.8}$$

All of the factors in the summand are constant when averaged over edge costs, with the exception of $\mathbb{I}[G \leq p]$. Referring back to the definition (3.4), we perform the average over edge costs to obtain the graphical expansion

$$C_{\mathbf{x},\mathbf{y}}(\ell_0) = \sum_{G \in \mathcal{G}_{\mathbf{x},\mathbf{y}}} d(G) \Pr[G \leq \ell_0], \tag{3.9}$$

where I introduce a diagrammatic quantity

$$d(G \in \mathcal{G}_{\mathbf{x},\mathbf{y}}) \equiv \sum_{\Gamma' \subseteq \Gamma_{\mathbf{x},\mathbf{y}}(G)} (-1)^{|\Gamma'|+1} \mathbb{I}[\Gamma' \text{ covers } G], \tag{3.10}$$

which is independent of the edge costs and referred to by Essam as the d-weight of the graph G.

Alternate form for the d-weight

The definition of $d(G)$ may be extended to cover the case where G is any two-rooted graph as follows: if G consists of more than one connected component, there is no way to cover all its edges with paths connecting the roots, so $d(G) = 0$. Note that, because the "covering" criterion is defined in terms of the edge set only, addition of isolated vertices does not change a graph's d-weight. Similarly, if G is not vertex-irreducible, by definition some edges — the "tadpoles" or "dangling ends" — cannot be covered by a self-avoiding path, since backtracking is forbidden, so again $d(G) = 0$. Since $d(G)$ vanishes for these additional cases, the sum in (3.9) may be extended to all two-rooted subgraphs of the underlying lattice.

The preceding derivation also applies to connectedness functions on an arbitrary graph G instead of the whole lattice; the sum in (3.10) is then over appropriate subgraphs of G. I define this generalized function by $C_{\mathbf{x},\mathbf{y}}(G,p)$. Equation (3.9) generalizes to

$$C_{\mathbf{x},\mathbf{y}}(G,p) = \sum_{E' \subseteq E(G)} d(G'_{E'}) \Pr[G_{E'} < p], \tag{3.11}$$

where $G'_{E'}$ is the subgraph of G consisting of all G's vertices and a subset $E' \subseteq E$ of its edges. Evaluating (3.11) at $p = 1$ yields

$$\mathbb{I}[E(G) \text{ connects } \mathbf{x}, \mathbf{y}] = \sum_{E' \subseteq E(G)} d(G'_{E'}). \tag{3.12}$$

Since G is now an arbitrary graph and I have extended the sum in (3.11) to cover all subsets of $E(G)$, we may easily invert this sum by Möbius inversion [93], which for this case merely reproduces inclusion-exclusion. We obtain

$$d(G) = \sum_{E' \subseteq E(G)} (-1)^{|E(G)|-|E'|} \mathbb{I}[E' \text{ connects } \mathbf{x}, \mathbf{y}], \tag{3.13}$$

which is equivalent to (3.10) but easier to work with in later proofs in that the sum is more straight-

forward than the one over self-avoiding paths. It is worth remarking that the sum over paths was only used to arrive at the form (3.9) for arbitrary G. Once this has been done, the coefficients (3.13) were obtained with no further reference to paths.

3.1.2 n-point connectedness functions

The expansion (3.8) generalizes readily to n-point connectivity functions [94, 95]; the criterion is simply that n root points $\mathbf{x}_1, \cdots, \mathbf{x}_n$ are connected if and only if there exists at least one path from $\mathbf{x}_1$ to each $\mathbf{x}_i$, $2 \geq i \geq n$, where we select $\mathbf{x}_1$ arbitrarily. Note that in enumerating the set of paths from x_1 to x_i, we must include those paths that pass through other root points. Using inclusion-exclusion (3.2) and equation (3.5) again, we may write the indicator function for this event as

$$\begin{aligned} \mathbb{I}[\mathbf{x}_1, \cdots, \mathbf{x}_n \text{ connected at } p] &= \prod_{i=2}^{n} \mathbb{I}_c(\mathbf{x}_1, \mathbf{x}_i)|_p & (3.14) \\ &= \prod_{i=2}^{n} \sum_{\emptyset \subset \Gamma_i' \subset \Gamma_{\mathbf{x}_1, \mathbf{x}_i}} (-1)^{|\Gamma_i'|+1} \mathbb{I}\Big[\bigwedge_{\gamma \in \Gamma_i'} (\gamma \leq p) \Big]. & (3.15) \end{aligned}$$

Repeating the previous derivation and grouping together terms that test the same set of edges, we obtain the diagrammatic expansion

$$C_{\mathbf{x}_1, \cdots, \mathbf{x}_n}(p) = \sum_{G \in G_{\mathbf{x}_1, \cdots, \mathbf{x}_n}} d(G) \Pr[G \leq p], \tag{3.16}$$

where the n-point d-weight is

$$d(G \in G_{\mathbf{x}_1, \cdots, \mathbf{x}_n}) \equiv \prod_{i=2}^{n} \sum_{\emptyset \subset \Gamma_i' \subset \Gamma_{\mathbf{x}_1, \mathbf{x}_i}(G)} (-1)^{|\Gamma_i'|+1} \mathbb{I}[\cup_i \Gamma_i' \text{ covers } G]. \tag{3.17}$$

The argument following (3.10) also carries though, since the above definition of the d-weight may be extended to arbitrary graphs and we may perform Möbius inversion on the connectedness

function evaluated on an arbitrary n-point graph, obtaining

$$d(G \in G_{\mathbf{x}_1,\cdots\mathbf{x}_n}) = \sum_{E' \subseteq E(G)} (-1)^{|E(G)|-|E'|} \mathbb{I}\left[E' \text{ connects } \mathbf{x}_1, \cdots \mathbf{x}_n\right]. \tag{3.18}$$

Our final results, equations (3.16) and (3.18), constitute a complete low-density epxansion for all connectedness properties of percolation clusters. I now show that this is entirely equivalent to the conventional treatment in terms of the $Q \to 1$ limit of the Q-state Potts model.

3.1.3 Equivalence with the Potts model

Potts model partition function

The development of the field theory for the Potts model is described in detail elsewhere [89, 90, 96, 97] and I recall only the parts of the derivation which are relevant to my discussion here. The Q-state Potts model on a graph G [88] has a degree of freedom $\alpha(\mathbf{x})$ associated with each vertex $\mathbf{x}$ of G which may take on any of Q discrete values (usually referred to as colors). The Hamiltonian for this model in the absence of a source field is

$$H_{\text{Potts}} = -J \sum_{\langle \mathbf{x},\mathbf{x}' \rangle} (\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} - 1), \tag{3.19}$$

where the sum is over pairs of neighboring vertices $\mathbf{x}, \mathbf{x}'$. The partition function can be expanded in the form [98]

$$Z = \sum_{\{\alpha(\mathbf{x})\}} e^{-\beta H_{\text{Potts}}} = \sum_{\{\alpha\}} \prod_{\langle \mathbf{x},\mathbf{x}' \rangle} \left[(1 - e^{-\beta J})\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} + e^{-\beta J}\right] \tag{3.20}$$

$$= \sum_{E' \subseteq E} p^{|E'|}(1-p)^{|E|-|E'|} Q^{N_c(G_{E'})}. \tag{3.21}$$

In the second line I have expressed the partition function as a sum over cluster configurations $G_{E'}$ where an edge on E' is present if the two vertices it is incident upon have the same color. Here $p = 1 - e^{-\beta J}$ and $N_c(G_{E'})$ is the number of connected components of the graph $G_{E'}$. Viewed as a function of Q and p, this sum is closely related to the Tutte polynomial of the graph G [91].

If we take $Q \to 1$, the above partition function corresponds to the sum of probabilities for the sets E' of occupied edges in bond percolation with independent probabilities p for occupying each edge. However, exactly at $Q = 1$ the partition function is trivially $Z = 1$, so in order to use the partition function as a generating functional for correlation functions we must work at arbitrary Q and then *a posteriori* evaluate the derivative with respect to Q of the resulting expressions at $Q = 1$ [90]. (Note that, on the other hand, the critical exponents of percolation are those of the Q-state Potts model evaluated exactly at $Q = 1$). This limitation (that one cannot directly work with the correlation functions of percolation from the start) necessitates my use of Essam's expansion.

The internal symmetry group of the Potts model is S_Q, the permutation group on Q objects. This may be made manifest by letting the lattice degrees of freedom consist of an overcomplete set of Q vectors $\vec{e}^{\alpha}$, $\alpha = 1, \ldots, Q$, in a $Q-1$ dimensional space. These vectors are obtained by projecting the position vectors of a regular Q-simplex in Q-dimensional space onto the subspace orthogonal to the vector $(1, 1, \ldots, 1)$. More concretely, if we let the coordinates of these vectors with respect to some basis be e_i^{α}, $i = 1, \ldots, Q-1$, the set of vectors may be uniquely defined up to relabeling and change of basis by requiring that

$$\sum_{\alpha=1}^{Q} e_i^{\alpha} = 0, \tag{3.22}$$

$$\sum_{\alpha=1}^{Q} e_i^{\alpha} e_j^{\alpha} = Q\delta_{ij}, \text{ and} \tag{3.23}$$

$$\sum_{i=1}^{Q-1} e_i^{\alpha} e_i^{\beta} = Q\delta^{\alpha\beta} - 1. \tag{3.24}$$

In equations (3.22) – (3.24), I have normalized the vectors following the convention used in [90, 96, 99]. Note that [89] and [97] adopt a different normalization.

By (3.24) the Potts interaction may be rewritten using

$$\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} = \frac{1}{Q}\left(\sum_{i=1}^{Q-1} e_i^{\alpha(\mathbf{x})} e_i^{\alpha(\mathbf{x}')} - 1\right); \tag{3.25}$$

we may think of the variables $\{i\}$ being associated with the edges of G just as the $\{\alpha\}s$ are asso-

ciated with its vertices.

Potts model correlation functions

To obtain the two-point connectedness function, we introduce factors $e_{i_1}^{\alpha(\mathbf{x}_1)}$, $e_{i_2}^{\alpha(\mathbf{x}_2)}$ into the partition sum. If they are not in the same connected component in the expansion, the sum over all α's gives zero by (3.22). That is,

$$C_{\mathbf{x}_1,\mathbf{x}_2}(p) = \frac{d}{dQ}\sum_{i_1,i_2} C_{i_1,i_2}(\mathbf{x}_1,\mathbf{x}_2)\Big|_{Q=1}. \tag{3.26}$$

where

$$C_{i_1,i_2}(\mathbf{x}_1,\mathbf{x}_2) \equiv \sum_{\{\alpha\}} e_{i_1}^{\alpha(\mathbf{x}_1)} e_{i_2}^{\alpha(\mathbf{x}_2)} \prod_{\langle \mathbf{x},\mathbf{x}'\rangle} \left[p\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} + (1-p)\right] \tag{3.27}$$

$$= \delta_{i_1,i_2} \sum_{E'\subseteq E} \mathbb{I}\left[E' \text{ connects } \mathbf{x}_1,\mathbf{x}_2\right] p^{|E'|}(1-p)^{|E|-|E'|} Q^{N_c(G_{E'})}, \tag{3.28}$$

where we also used (3.23).

Now we rewrite

$$p\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} + (1-p) = p(\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} - 1) + 1. \tag{3.29}$$

(Although this does not explicitly involve the e_i^α's, this choice is motivated by the form of eq. (3.24) as $Q \to 1$; note that there are many similar expressions that become equal to this for $Q = 1$.) We expand the Potts correlation function $C_{i_1,i_2}(\mathbf{x}_1,\mathbf{x}_2)$ using this decomposition for each edge, and then once more for $\delta_{\alpha(\mathbf{x}),\alpha(\mathbf{x}')} - 1$ on each edge. This yields

$$C_{i_1,i_2}(\mathbf{x}_1,\mathbf{x}_2) = \delta_{i_1,i_2} \sum_{E'\subseteq E(G)} p^{|E'|} d_Q(G_{E'}), \tag{3.30}$$

where

$$d_Q(G_{E'}) = \sum_{E''\subseteq E'} (-1)^{|E'|-|E''|} \mathbb{I}\left[E'' \text{ connects } \mathbf{x}_1,\mathbf{x}_2\right] Q^{N_c(E'')}. \tag{3.31}$$

We pick up a factor of $Q-1$ from $\sum_{i_1,i_2}\delta_{i_1,i_2}$; removing this and setting $Q=1$ is equivalent to evaluating the derivative in (3.26). We therefore recover the expressions (3.11) and (3.13). The derivation can be readily generalized, at least to the 3-point connectedness function, so I have proved Essam's diagrammatic expansion is identical term-by-term with the low-density expansion of the Potts model in the $Q \to 1$ limit.

3.1.4 Topological properties of the diagrammatic weights

One implication of the results of the preceding sections is that the d-weight is a "topological invariant" of the graph G in the sense that $d(G)$ is unchanged by replacing edges of G with chains of edges (*i.e.*, inserting vertices of degree two). This means that we can consider the expansion in terms of what I call topological graphs $\mathcal{G} \in \mathcal{G}_{\mathbf{x},\mathbf{y}}$, equivalence classes of lattice graphs obtained by removing any degree-two vertices. These are simply graphs, without any embedding in the lattice, with two distinguished vertices labeled $\mathbf{x}, \mathbf{y}$. Using this property, we may rewrite the expansion (3.9) as

$$C_{\mathbf{x},\mathbf{y}}(p) = \sum_{\mathcal{G}\in\mathcal{G}_{\mathbf{x},\mathbf{y}}} \frac{d(\mathcal{G})}{\mathcal{A}(\mathcal{G})} \sum_{\lambda:\mathcal{G}\to G} \Pr[G \le p], \tag{3.32}$$

with analogous expressions for the n-point functions. Here the second sum is over all possible embeddings $\lambda : \mathcal{G} \to G$ which map the edges of $\mathcal{G}$ into self-avoiding chains of edges on the lattice, producing the set of lattice graphs summed over in (3.9). Note that these chains must not only be self-avoiding walks, but also must avoid intersection with chains arising from different edges of $\mathcal{G}$.

If the topological graph $\mathcal{G}$ has any non-trivial automorphisms which leave the root points $\mathbf{x}$, $\mathbf{y}$ fixed, then there is more than one embedding λ which produces the same lattice graph G. Consequently, we must divide by the number $\mathcal{A}(\mathcal{G})$, which is the number of elements in the automorphism group of $\mathcal{G}$. This number is known as the "symmetry factor" from the theory of Feynman diagrams (see *e.g.*[100–102]), and will play such a role in the continuum theory.

This topological property of the d-weights is crucial for extending the lattice expansion (3.32) to a continuum theory, a point which I will return to in section 3.3.1. It also simplifies lattice calculations, since it greatly reduces the number of different graphs for which $d(G)$ must be calculated.

Singly-connected edges

I now introduce another quantity in the lattice expansion which will have a useful counterpart in the theory of MSF paths. This is the derivative of $C_{\mathbf{x},\mathbf{y}}(p)$ with respect to the value p_e of p on a particular edge e: the probability that $\mathbf{x}$ and $\mathbf{y}$ are connected by a cluster at p changes at $p_e = p$ only if $\mathbf{x}$ and $\mathbf{y}$ are not connected when $p_e = p - dp$, and are connected when $p_e = p$. This implies that for $p_e = p$, any path from $\mathbf{x}$ to $\mathbf{y}$ on the cluster must traverse e [103]. An edge with this property is called a singly-connected edge, also known as "red bonds" in the terminology of [39, 104]. We may define

$$C^e_{\mathbf{x},\mathbf{y}}(p) = \int_0^p dp_e \frac{\partial}{\partial p_e} C_{\mathbf{x},\mathbf{y}}(p, p_e), \tag{3.33}$$

which is the probability that at parameter p, $\mathbf{x}$ and $\mathbf{y}$ are connected, and e is a singly-connected edge on the same cluster. Note that this is not the same as the 3-point connectedness function that was defined above, and its lattice expansion (which may be obtained straightforwardly, as a consequence of the definition, in terms of that for $C_{\mathbf{x},\mathbf{y}}(p)$) still contains the same d weights as for $C_{\mathbf{x},\mathbf{y}}(p)$. In fact, the entire expansion is remains the same as (3.32), with the additional condition that the edge e belong to all embedded graphs G in the series. Note that due to the formulation of this expansion via inclusion-exclusion, the edge e does *not* have to be singly connected in each term.

In terms of topological graphs, we may view the series as extending over those graphs having the inverse image of e as a marked edge (either of the ends of which may be degree-two vertices), which is mapped to the single edge e under all embeddings λ. In this case the relevant automorphisms of $\mathcal{G}$ must fix this edge as well as $\mathbf{x}$, $\mathbf{y}$, and I denote the number of these by $\mathcal{A}'(\mathcal{G})$. Clearly $\mathcal{A}'(\mathcal{G}) \leq \mathcal{A}(\mathcal{G})$: the one group of automorphisms is a subgroup of the other. These different ways of writing the function are equivalent.

I now proceed to develop a graphical expansion for MSF connectedness functions, analogous to (3.32) in that it takes the form of a weighted sum over topological graphs and their lattice embeddings. This will be the basis for the continuum theory analyzed in section 3.3.

3.2 A graphical expansion for MSF paths

In this section, I define a lattice expansion for the equivalent of the "one-point" connectedness function $P_{(0;m)}$ on the Bethe lattice: this is the probability that two points $\mathbf{x}, \mathbf{y}$ are connected by MSF(p_0) *and* the connecting path passes through a third point $\mathbf{z}$, which I will refer to as the "path vertex." I denote this object by $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p_0)$, and its properties will yield the fractal dimension of MSF paths D_p. If we were to pursue the analogy with percolation as closely as possible, it would be more appropriate to work with the probability the MSF path passes through an edge e instead of a vertex $\mathbf{z}$ (by analogy with (3.33)); however, these quantities are equivalent in the continuum limit.

I obtain a diagrammatic expansion for $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p_0)$ by imposing the no-cycle condition involved in the Kruskal process to the connectedness expansion for bond percolation obtained in the previous section. Again, for simplicity, I restrict my discussion to the situation where only finite cycles are formed, which means that we must take $p_0 \leq p_c$, and hence the connectedness functions I calculate are those for the MSF constructed on finite percolation clusters.

An important difference between percolation clusters and MSF(p_0) is that any two points may be connected by at most one path on the latter, namely the MST path with those endpoints if it has formed by p_0. In section 3.2.2 I develop a method to identify this path, based on properties of the MST which are first proved in section 3.2.1. I then formulate a diagrammatic expansion for $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p_0)$ in terms of lattice graphs in section 3.2.3 and in terms of topological graphs in section 3.2.4. My principal results are the expansions (3.50), (3.67) with new diagrammatic weights given in (3.49), (3.54).

3.2.1 MST paths as geodesics

We identify MSF paths using their *geodesic* property, briefly alluded to in section 2.1.1: the path $\gamma_{\text{MST}}(\mathbf{x}, \mathbf{y})$ on the completed MST connecting two points $\mathbf{x}, \mathbf{y}$ is the unique such path which minimizes the maximum edge cost for *all* contiguous subsets of the path, which I now show.

Definitions

In what follows, we examine the set of all self-avoiding paths with fixed endpoints $\mathbf{x}, \mathbf{y}$ on some arbitrary finite graph G, which I denote by $\Gamma_{\mathbf{x},\mathbf{y}}$. A path $\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}$ is a minimax path if, out of all paths in $\Gamma_{\mathbf{x},\mathbf{y}}$, it minimizes the cost of the most expensive edge (which I refer to as the barrier) on the path; *i.e.*it is a minimum over $\Gamma_{\mathbf{x},\mathbf{y}}$ of the maximum (over edges on a path) of the cost of an edge. In general, there are many minimax paths between two given points, since that property only concerns a single edge out of all the edges on the path. This point does not appear to be well appreciated in the literature, *e.g.* [10, 12, 20], which will frequently refer in the singular to "the" minimax path between two points.

These sources really mean the geodesic path, which is defined as a path $\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}$ such that, for any vertices $\mathbf{w}, \mathbf{z} \in \gamma$, the subset $\gamma' \subseteq \gamma$ connecting these vertices is a minimax path from the set $\Gamma_{\mathbf{w},\mathbf{z}}$. In particular, this holds for cases where $\mathbf{w}, \mathbf{z}$ are adjacent vertices connected by a single edge, so we may equivalently define geodesics as those paths whose edges are all minimax paths connecting the vertices to which they are incident. This implies that geodesics cannot be cycles.

I empahsize that the geodesic path is the correct strong disorder limit of the optimal path [19], defined as the path minimizing the total cost of all edges on the path; hence these paths are of central importance for strong disorder transport problems.

Existence and uniqueness

I now prove that the geodesic between arbitrary endpoints $\mathbf{x}, \mathbf{y}$ is unique, if it exists. Specifically, I let γ be a geodesic and show that no other path $\overline{\gamma} \neq \gamma$ which shares the same endpoints may also be a geodesic. Clearly $\overline{\gamma}$ and γ must both be minimax paths and so must pass through the same most costly edge e_1, possibly in different directions. Label the incident vertices of e_1 as $\mathbf{w}_1, \mathbf{z}_1$ such that γ traverses the points in the order $\mathbf{x}, \mathbf{w}_1, \mathbf{z}_1, \mathbf{y}$. It then follows that any minimax path from $\mathbf{x}$ to $\mathbf{w}_1$ must pass through a distinct barrier edge e_2, whose cost is less than that of e_1 by the assumption of distinct edge costs. This means that no path transversing the four vertices in the order $\mathbf{x}, \mathbf{z}_1, \mathbf{w}_1, \mathbf{y}$ can be a minimax path on the segment from $\mathbf{x}$ to $\mathbf{w}_1$, since the barrier edge of such a segment would be e_1. If $\mathbf{w} = \mathbf{x}$, then we can start from $\mathbf{y}$ instead, and if e_1 has endpoints

$\mathbf{x}$ and $\mathbf{y}$ then we are done.

This establishes that $\overline{\gamma}$ transverses the four identified points in the same order as γ. By repeatedly splitting up γ at the barrier edges of its subpaths, we establish that $\overline{\gamma}$ must pass through the same vertices as γ in the same order, hence the geodesic γ is unique.

As a consequence of the minimax path property (section 2.1.1), all paths on the MST are geodesics. I now prove that all geodesics are MST paths: if the edge $e_{i,j}$ connecting adjacent vertices i, j is also the minimax path between i and j, it is by definition the cheapest edge across the cut constructed by taking all paths from i to j and removing the most expensive edge on each. By the strong cut property (section 2.1.1), the existence of such a cut means that the edge $e_{i,j}$ is on the MST, and as noted above, all geodesics are a union of barrier edges like $e_{i,j}$. Therefore, the MST of a graph is the set of all of its barrier edges. Since a MST path exists for arbitrary endpoints, provided they lie on the same connected component of G, this establishes the existence of a geodesic under the same condition.

Ultrametricity

It's worth mentioning that the minimax property of the MST gives rise to an (ultra)metric structure from which the geodesic paths may be obtained — this justifies my choice of the term "geodesic" for these paths.

For the remainder of this section, I assume the costs are non-negative, which can be guaranteed without loss of generality by shifting them by a positive constant. The cost of the barrier edge then gives rise to a distance measure ∂, or metric, between any vertices $\mathbf{x}, \mathbf{y}$ of G. As noted above, this cost is uniquely defined for $\mathbf{x} \neq \mathbf{y}$; I define $\partial(\mathbf{x}, \mathbf{y}) = 0$. By definition, a metric should be finite and non-negative, symmetric ($\partial(\mathbf{x}, \mathbf{y}) = \partial(\mathbf{y}, \mathbf{x})$), equal to zero if and only if $\mathbf{x} = \mathbf{y}$, and obey the triangle inequality. The first three properties are clear from properties of minimax paths given above. In fact, ∂ obeys a stronger form of the triangle inequality: for any $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$,

$$\partial(\mathbf{x}, \mathbf{y}) \leq \max\left(\partial(\mathbf{x}, \mathbf{z}), \partial(\mathbf{z}, \mathbf{y})\right). \tag{3.34}$$

These four properties imply that ∂ is an ultrametric. For ordinary metric spaces, one defines

geodesics to be paths of shortest "length" using the metric, and this motivates my terminology above. Further, if $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$ are three distinct points, the ultrametric property implies that if $\partial(\mathbf{x},\mathbf{y}) \leq \partial(\mathbf{x},\mathbf{z})$ and $\partial(\mathbf{y},\mathbf{z})$, then $\partial(\mathbf{x},\mathbf{z}) = \partial(\mathbf{y},\mathbf{z})$.

It is well known that an ultrametric space with a finite number of points can be viewed as a tree, which we imagine depicted with the points as the leaves located on a hyperplane, other vertices to one side of the hyperplane, connected by straight lines, and the ultrametric represented by the height above (in the direction orthogonal to the hyperplane) the leaves to which one must go in walking from one leaf to another along the tree. In the present case, this essentially corresponds to the MST. The tree is trivalent (except at the leaves) with probability one. The trivalent vertices represent the edges on the MST, with their height as their cost. In fact, if we consider the subforest of the tree consisting of the vertices at height less than or equal to some bound, then this gives MSF(p).

3.2.2 Identifying MST paths through binary comparisons

I now denote the geodesic path between $\mathbf{x},\mathbf{y}$ as $\gamma_{\text{MST}}(\mathbf{x},\mathbf{y})$. The geodesic properties derived above allow $\gamma_{\text{MST}}(\mathbf{x},\mathbf{y})$ to be selected from the set $\Gamma_{\mathbf{x},\mathbf{y}}$ of all paths connecting $\mathbf{x},\mathbf{y}$ by means of repeated comparisons using a binary ordering relation $\prec$, defined as follows. Let γ and γ' be two paths in $\Gamma_{\mathbf{x},\mathbf{y}}$. Let $e_1, e_2, \cdots e_n$ be the first, second, ... n-th most expensive edges on γ, and likewise for $e'_1, e'_2, \cdots e'_n$ on γ'. We say $\gamma \prec \gamma'$ if and only if there exists some j such that $\ell_{e_j} < \ell_{e'_j}$ and $\ell_{e_i} = \ell_{e'_i}$ for all $i > j$: in other words, we compare the most expensive edges whose costs are not identical. I now prove that

$$\gamma_{\text{MST}}(\mathbf{x},\mathbf{y}) = \min_{\gamma\in\Gamma_{\mathbf{x},\mathbf{y}}} \gamma; \tag{3.35}$$

in other words, $\gamma_{\text{MST}}(\mathbf{x},\mathbf{y})$ is the minimal element of the set $\Gamma_{\mathbf{x},\mathbf{y}}$ under the ordering defined by $\prec$.

I assume all edge costs to be distinct (which, of course, holds with probability 1 in the case where the ℓ_e are chosen from a continuous distribution), which implies that $e_i = e'_i \iff \ell_{e_i} = \ell_{e'_i}$ and $\neg(\gamma \preceq \gamma') \iff \gamma \succ \gamma'$: *i.e.*, the relation $\prec$ defines a *total* order on the set of all paths between

fixed endpoints. Under this assumption, for any two paths $\gamma, \overline{\gamma} \in \Gamma_{\mathbf{x},\mathbf{y}}$ we have

$$\gamma \prec \overline{\gamma} \iff \max_{e \in \gamma - \gamma \cap \overline{\gamma}} \ell_e < \max_{e' \in \overline{\gamma} - \gamma \cap \overline{\gamma}} \ell_{e'}. \tag{3.36}$$

We let γ satisfy

$$\gamma = \min_{\overline{\gamma} \in \Gamma_{\mathbf{x},\mathbf{y}}} \overline{\gamma}. \tag{3.37}$$

In particular, γ is less than all elements in $\Gamma(\gamma; \mathbf{x}, \mathbf{y}) \subset \Gamma_{\mathbf{x},\mathbf{y}}$, which I define as the set of paths in $\Gamma_{\mathbf{x},\mathbf{y}}$ having no edges in common with γ. By (3.36), this means $\prec$ compares only the most expensive edges on the paths, so if $\gamma \prec \overline{\gamma} \in \Gamma(\gamma; \mathbf{x}, \mathbf{y})$ then γ must be a minimax path.

Similarly, for any subpath $\gamma' \subseteq \gamma$ with endpoints $\mathbf{w}$, $\mathbf{z}$, I define $\Gamma(\gamma; \mathbf{w}, \mathbf{z}) \subset \Gamma_{\mathbf{x},\mathbf{y}}$ as the set of those paths that differ from γ only on this subpath; in other words the set of those $\overline{\gamma}$ such that $\gamma - \gamma \cap \overline{\gamma} = \gamma'$. If $\gamma \prec \overline{\gamma} \in \Gamma(\gamma; \mathbf{w}, \mathbf{z})$, then γ' is a minimax path from $\mathbf{w}$ to $\mathbf{z}$, and since our assumption implies that this holds for all choices of $\mathbf{w}$ and $\mathbf{z}$, this means that γ is a geodesic, and hence the unique path on the MST by the results of the previous section.

Remarks

Mathematically inclined readers will notice that I have been cavalier in assuming the existence of a minimal γ in the continuum limit, where the set $\Gamma_{\mathbf{x},\mathbf{y}}$ becomes infinite. This is justifiable for my purposes, since I may construct the continuum limit as a sequence of increasingly large finite lattices. In section 2.2.3, I defined $\mathbf{x}$ and $\mathbf{y}$ to be connected only if $\gamma_{\text{MST}}(\mathbf{x}, \mathbf{y})$ does not intersect the boundary, so we are only interested in finding bounded paths.

Furthermore, in this chapter we only examine the case $p_0 \leq p_c$. By Kruskal's algorithm (section 2.1.1), the MSF at p_0 is supported on the percolation clusters at p_0, and the fact that we don't go above the percolation threshold means that the connected components of the MSF at p_0 containing $\mathbf{x}$ and $\mathbf{y}$ will be bounded with probability 1. This means that, despite my notation, we effectively always deal with a finite $\Gamma_{\mathbf{x},\mathbf{y}}$, which always contains its minimal and maximal elements.

I finally note that identifying MST paths through the procedure (3.37) is most convenient for the purposes of my diagrammatic expansion (3.67), despite being very computationally inefficient,

since we make many unnecessary comparisons with paths that aren't in any of the sets $\Gamma(\gamma; \mathbf{w}, \mathbf{z})$. On the other hand, the geodesic characterization of MST paths is less directly useful for my purposes, since it requires keeping track of the locations of the most expensive edges. The geodesic properties of MST paths are very useful computationally: they are essential to constructing linear-time algorithms for MST path verification [105] which were used in [106] to give a randomized algorithm which constructs the entire MST in linear time.

3.2.3 A lattice expansion for MSF paths

I now seek an expansion, analogous to (3.8), for the probability that the points $\mathbf{x}, \mathbf{y}$ are connected by a path on the MSF at p which passes through $\mathbf{z}$, which I denote by

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p) \equiv \langle \mathbb{I}[\mathbf{x}, \mathbf{y} \text{ connected at } p \textit{ and } \gamma_{\text{MST}}(\mathbf{x}, \mathbf{y}) \text{ passes through } \mathbf{z}] \rangle \tag{3.38}$$

$$= \langle \mathbb{I}_c(\mathbf{x}, \mathbf{y}; \mathbf{z})|_p \rangle. \tag{3.39}$$

Due to the greedy nature of Kruskal's algorithm, we never remove edges from the MSF as p is increased, so if a path connecting two points on the MSF at a parameter value p exists, it must be identical to the unique path connecting those points on the completed MST. In section 3.2.2 I used the geodesic properties of the MST path to arrive at the definition

$$\gamma_{\text{MST}}(\mathbf{x}, \mathbf{y}) = \min_{\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}} \gamma, \tag{3.40}$$

where '$\min$' denotes the minimal element under the relation $\prec$ defined in (3.36). This lets us write the indicator function in (3.38) as

$$\mathbb{I}_c(\mathbf{x}, \mathbf{y}; \mathbf{z})|_p = \mathbb{I}\big[\min_{\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}} \gamma \leq p\big] \mathbb{I}_c^{(\mathbf{z})}\big[\min_{\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}} \gamma\big], \tag{3.41}$$

where for brevity, I introduce

$$\mathbb{I}_c^{(\mathbf{z})}[\gamma] \equiv \mathbb{I}[\gamma \text{ passes through } \mathbf{z}]. \tag{3.42}$$

and, again, $\Gamma_{\mathbf{x},\mathbf{y}}$ is the set of all self-avoiding paths on G between $\mathbf{x}$ and $\mathbf{y}$.

In order to write this as a sum over all paths $\gamma \in \Gamma_{\mathbf{x},\mathbf{y}}$, I introduce a third indicator function to select the minimal path from the sum, which I write as

$$\mathbb{I}_c(\mathbf{x},\mathbf{y};\mathbf{z})|_p = \sum_{\gamma\in\Gamma_{\mathbf{x},\mathbf{y}}} \mathbb{I}[\gamma \le p]\mathbb{I}_c^{(\mathbf{z})}[\gamma]\mathbb{I}\Big[\bigwedge_{\gamma'\in\Gamma_{\mathbf{x},\mathbf{y}}} (\gamma \preceq \gamma')\Big]. \tag{3.43}$$

Expanding this third indicator function by inclusion-exclusion gives

$$\mathbb{I}_c(\mathbf{x},\mathbf{y};\mathbf{z})|_p = \sum_{\gamma\in\Gamma_{\mathbf{x},\mathbf{y}}} \mathbb{I}[\gamma \le p]\mathbb{I}_c^{(\mathbf{z})}[\gamma] \times \sum_{\Gamma'\subseteq\Gamma_{\mathbf{x},\mathbf{y}}} (-1)^{|\Gamma'|}\mathbb{I}\Big[\bigwedge_{\gamma'\in\Gamma'} \neg(\gamma \preceq \gamma')\Big]. \tag{3.44}$$

Because the uniqueness of the edge costs implies $\neg(\gamma \preceq \gamma') \iff (\gamma \succ \gamma')$, we may restrict the sum over subsets of $\Gamma_{\mathbf{x},\mathbf{y}}$ to those not containing γ itself. We now reorganize the double sum by grouping together all terms that test the same set of edges, as was done for equation (3.8) for percolation. For each term in (3.44), the edges in $\gamma \cup \Gamma'$ will form a graph $G \in \mathcal{G}_{\mathbf{x},\mathbf{y}}$, the set of all vertex-irreducible graphs with root vertices $\mathbf{x}, \mathbf{y}$. When we regroup the sum in terms of these graphs, we will obtain a sum over sets $\Gamma'' = \gamma \cup \Gamma'$ of paths from $\mathbf{x}$ to $\mathbf{y}$ which cover G and contain the chosen path γ, similar to what was obtained to percolation. Unlike the percolation case, we still have the outermost sum in (3.44), which will now run over elements of Γ'', so the graphical expansion is

$$\mathbb{I}_c(\mathbf{x},\mathbf{y};\mathbf{z})|_p = \sum_{G\in\mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}} \mathbb{I}[G \le p] \sum_{\Gamma''\subseteq\Gamma_{\mathbf{x},\mathbf{y}}(G)} \mathbb{I}[\Gamma'' \text{ covers } G] \sum_{\gamma\in\Gamma''} (-1)^{|\Gamma''|+1}\mathbb{I}\Big[\bigwedge_{\gamma'\in\Gamma''-\gamma} (\gamma \succ \gamma')\Big]\mathbb{I}_c^{(\mathbf{z})}[\gamma]. \tag{3.45}$$

As in the derivation of (3.9), we may factor out the dependence on the parameter p as

$$\mathbb{I}_c(\mathbf{x},\mathbf{y};\mathbf{z})|_p = \sum_{G\in\mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}} d_{\text{MSF}}(G)\mathbb{I}[G \le p], \tag{3.46}$$

where I have introduced $d_{\text{MSF}}(G)$, the analogue of Essam's d-weight (3.10) for MSF paths:

$$d_{\text{MSF}}(G) \equiv \sum_{\Gamma' \subseteq \Gamma_{\mathbf{x},\mathbf{y}}(G)} (-1)^{|\Gamma'|+1} \mathbb{I}[\Gamma' \text{ covers } G] \sum_{\gamma \in \Gamma'} \mathbb{I}\Big[\bigwedge_{\gamma' \in \Gamma' - \gamma} (\gamma \succ \gamma') \Big] \mathbb{I}_c^{(\mathbf{z})}[\gamma]. \tag{3.47}$$

In (3.47), I have suppressed the dependence of $d_{\text{MSF}}(G)$ on the points $\mathbf{x}, \mathbf{y}, \mathbf{z}$ and on the edge costs of G.

In the analogous statement (3.10) for percolation, we found $d(G)$ was independent of edge costs and depended on topological properties of G only. Here we are not so lucky: to evaluate $d_{\text{MSF}}(G)$ we need to be able to compare paths which cover G using the relation $\succ$. By the definition (3.36), a necessary and sufficient set of information to do this is the relative ordering of the edge costs of G. I incorporate this by indexing the edges of G arbitrarily and defining an ordering of their costs to be given by a permutation $\pi \in S_{|E(G)|}$ on the set of $|E(G)|$ elements, via

$$\ell_{\pi(i)} < \ell_{\pi(j)} \iff i < j. \tag{3.48}$$

With this notation, we see that d_{MSF} is a function of the graph G and edge cost ordering $\pi_{E(G)}$, so I write $d_{\text{MSF}}(G|\pi)$ (leaving the dependence on $\mathbf{x}, \mathbf{y}, \mathbf{z}$ implicit).

The second sum in (3.47) merely detects whether the maximal path in Γ' passes through the point $\mathbf{z}$, so for a fixed edge cost ordering π we may write

$$d_{\text{MSF}}(G|\pi) = \sum_{\Gamma' \subseteq \Gamma_{\mathbf{x},\mathbf{y}}(G)} (-1)^{|\Gamma'|+1} \mathbb{I}[\Gamma' \text{ covers } G] \mathbb{I}_c^{(\mathbf{z})}[\max_{\gamma \in \Gamma'} \gamma]. \tag{3.49}$$

Note that, as a consequence of our use of inclusion-exclusion, this definition is mildly counterintuitive: we are attempting to calculate the probability that the MSF path passes through $\mathbf{z}$, and by definition the MSF path (if it exists at p) is the minimum out of all paths in $\Gamma_{\mathbf{x},\mathbf{y}}$. However, for each graph in the perturbative expansion of this probability, the relevant event is that the *maximal* path of the covering Γ' passes through $\mathbf{z}$.

We may now take the expectation value of $\mathbb{I}_c(\mathbf{x}, \mathbf{y}; \mathbf{z})|_p$ over all realizations of the edge costs

in order to obtain the analogue of (3.9),

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p) = \sum_{G\in\mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}} \sum_{\pi_{E(G)}\in S_{|E(G)|}} d_{\text{MSF}}(G|\pi_{E(G)}) \Pr[(G \le p) \wedge \pi_{E(G)}]. \tag{3.50}$$

The ordering $\pi_{E(G)}$ and the event that all edges in G have cost less than p are independent, so the last probability in (3.50) factorizes. Because G is a graph embedded on the lattice and the edge costs are iid, all orderings of the edge costs are equally probable and $\Pr[\pi] = 1/|E(G)|!$; this will need to be modified modified when the expansion is done in terms of topological graphs in section 3.2.4.

Alternate form of d_{MSF}

We may also compute d_{MSF} in terms of edge subsets instead of covering paths, analogous to my derivation of (3.13) from (3.10). The argument proceeds the same way: we first generalize to MSF path connectedness functions on an arbitrary graph $G \in \mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$, denoted by $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(G,p)$. Because (3.49) contains a factor of $\mathbb{I}[\Gamma' \text{ covers } G]$ in the summand, we also have $d_{\text{MSF}}(G|\pi) = 0$ for disconnected or vertex-reducible graphs, and the sum in the generalization of (3.50) may be extended to all subgraphs of G as

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(G,p) = \sum_{\pi\in S_{|E(G)|}} \sum_{E'\subseteq E(G)} d_{\text{MSF}}(G_{E'}|\pi'_{E'}) \Pr[G_{E'} \le p] \Pr[\pi_{E'}], \tag{3.51}$$

where $\pi_{E'}$ is the relative ordering of the costs of the edges in E' induced by the ordering π_E. Because we know $d_{\text{MSF}}(G|\pi)$ is dependent on π, we must work under the sum over edge cost orderings in performing the Möbius inversion step. We therefore work with the conditional quantity

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(G,p|\pi) \equiv \sum_{E'\subseteq E(G)} d_{\text{MSF}}(G_{E'}|\pi'_{E'}) \Pr[G_{E'} \le p] \tag{3.52}$$

appearing as a summand in (3.51). Evaluating this at $p = 1$ yields

$$\mathbb{I}[E(G) \text{ connects } \mathbf{x},\mathbf{y}]\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\text{MST}}(G|\pi)] = \sum_{E'\subseteq E} d_{\text{MSF}}(G_{E'}|\pi'_{E'}), \tag{3.53}$$

where $\gamma_{\text{MST}}(G|\pi_{E(G)})$ is the path connecting the root points $\mathbf{x}, \mathbf{y}$ on the minimal spanning tree of G obtained under the edge cost ordering π. Möbius inversion of this sum gives

$$d_{\text{MSF}}(G|\pi) = \sum_{E' \subseteq E(G)} (-1)^{|E(G)|-|E'|} \mathbb{I}[E' \text{ connects } \mathbf{x}, \mathbf{y}] \mathbb{I}_c^{(\mathbf{z})}[\gamma_{\text{MST}}(G_{E'}|\pi'_{E'})]. \tag{3.54}$$

This definition of d_{MSF} is more convenient than (3.49) for the proofs of section 3.4.4. It is also, in principle, more convenient for computation, since for large graphs the size of the set $\Gamma_{\mathbf{x},\mathbf{y}}(G)$ of self-avoiding walks grows faster than $|E(G)|$, hence the sum in (3.54) is more easily performed than that in (3.49).

3.2.4 Expansion in terms of topological graphs

In the previous section, I gave a lattice expansion (3.50) for the path vertex connectedness function. If we wish to extend this expansion to a continuum theory, we must express this quantity as a sum over topological graphs $\mathcal{G}$. In particular, the combinatorics involved in enumerating the embeddings of a topological graph onto the lattice are highly non-trivial (see section 3.3.1 for further remarks), so in this section I seek to minimize the embedding dependence of the information required to compute (3.50).

I consider a given lattice graph G which corresponds to a topological graph $\mathcal{G}$ under some embedding λ. For clarity in what follows, I denote elements of the edge sets of $G, \mathcal{G}$ by different symbols: we have $e \in E(G)$ and $\epsilon \in E(\mathcal{G})$. I let $\lambda(\epsilon) \subset G$ denote the chain (self-avoiding path on the lattice) of N_ϵ edges on the lattice that the topological edge ϵ is mapped to.

Interestingly, for both percolation and the MST problem, the only edge cost information relevant for the connectedness functions is the cost of the most expensive edge on $\lambda(\epsilon)$, which I denote by

$$L_\epsilon = \max_{e \in \lambda(\epsilon)} \ell_e. \tag{3.55}$$

Then in percolation, to determine whether $\lambda(\epsilon)$ connects its endpoints, we only need to check whether $L_\epsilon \leq p$. Likewise, by the optimal path property, $\lambda(\epsilon)$ is on the MST of G only if L_ϵ is less than the maximal edge cost encountered on all other paths on $\mathcal{G}$ connecting the same endpoints.

In particular, since $d_{\text{MSF}}(G|\pi)$ is computed in terms of connectedness properties (whether the MST path between the root points of the diluted graph $G_{E'}$ goes through the MSF path vertex at $\mathbf{z}$), it may be computed using only the relative ordering of the $\{L_\epsilon\}$, which I denote by $\pi'_{E(\mathcal{G})}$. We may think of L_ϵ as an induced cost on the edge ϵ of $\mathcal{G}$, but its value depends on the embedding (through N_ϵ) and so $\Pr[\pi'_{E(\mathcal{G})}]$ does as well. Suppressing this dependence for now, we have

$$d_{\text{MSF}}(G|\pi) = d_{\text{MSF}}(\mathcal{G}|\pi'_{E(\mathcal{G})}). \tag{3.56}$$

and the lattice expansion (3.50) may be written as

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p) = \sum_{\mathcal{G}\in\mathcal{G}_{\mathbf{x},\mathbf{y},\mathbf{z}}} \sum_{\pi'\in S_{|E(\mathcal{G})|}} \frac{d_{\text{MSF}}(\mathcal{G}|\pi'_{E(\mathcal{G})})}{\mathcal{A}'(\mathcal{G})} \sum_{\lambda:\mathcal{G}\to G} \Pr[\pi'_{E(\mathcal{G})} \wedge (G \le p)], \tag{3.57}$$

by summing over all orderings of the non-maximal cost edges of G. Here $\mathcal{A}'(\mathcal{G})$ is the relevant symmetry factor as defined in section 3.1.4. Strictly speaking, it is defined here for automorphisms fixing the inverse image of the vertex $\mathbf{z}$ rather than an edge e; the cases relevant to the continuum expansion later are those in which the vertex $\mathbf{z}$ has degree two. The relevant point of comparison for the continuum theory is the case where e is shrunk to a single vertex of degree two. Eq. (3.57) is the main result of this section, and should be compared with the percolation result (3.32). In the following section I describe how the factor of $\Pr[\pi'_{E(\mathcal{G})} \wedge (G \le p)]$ may be put in a more convenient form — note that, as it stands, it may not be factorized as was done for the lattice expansion.

Assigning weights to the orderings

To put the quantity $\Pr[\pi'_{E(\mathcal{G})} \wedge (G \le p)]$ in a more tractable form, I return to basic considerations of Kruskal's algorithm on weighted lattice graphs. As the simplest example, I take the case where $\mathcal{G}$ consists of two root vertices connected in parallel by two edges ϵ_1, ϵ_2. I consider an embedding λ where $\lambda(\epsilon_1), \lambda(\epsilon_2)$ are chains of N_1, N_2 lattice edges, whose most expensive edge

has cost L_1, L_2, respectively. Obviously,

$$\Pr[G \leq \ell_0] = \prod_{e \in E(G)} p_0 = \prod_{\epsilon \in E(\mathcal{G})} p_0^{N_\epsilon}. \tag{3.58}$$

Note here I make use of the uniform rescaling of edge costs introduced in section 2.1.2. Again, the only difference between the percolation and Kruskal processes is that, in the latter, an edge with cost ℓ_e must not be accepted if it forms a cycle with edges accepted at $\ell_0 < \ell_e$. This means that we must modify the factor $\Pr[G \leq \ell_0]$ so that it keeps track of which of the paths corresponding to edges of $\mathcal{G}$ form first as ℓ_0 is raised. We may accomplish this by first generalizing (3.58) to the multivariate probability

$$\begin{aligned} P_G(\ell_1, \ell_2) &\equiv \Pr[(L_1 \leq \ell_1) \wedge (L_2 \leq \ell_2)] \\ &= p_1^{N_1} p_2^{N_2}. \end{aligned} \tag{3.59}$$

Clearly, we recover $\Pr[G \leq \ell_0]$ as $P_G(\ell_0, \ell_0)$. The utility of generalizing to P_G is that it allows us to calculate the probability of orderings on $\mathcal{G}$ through a process of differentiation and integration similar to that performed on the Bethe lattice. For example, the probability that G has formed by ℓ_0 and that its most expensive edge was on $\lambda(\epsilon_2)$ is

$$\Pr[L_1 < L_2 \leq \ell_0] = \int_{-\infty}^{\ell_0} d\ell_2 \int_{-\infty}^{\ell_2} d\ell_1 \, \frac{d}{d\ell_1} \frac{d}{d\ell_2} P_G(\ell_1, \ell_2) \tag{3.60}$$

$$= \frac{N_2}{N_1 + N_2} p_0^{N_1 + N_2}. \tag{3.61}$$

This is self-evident from the iid nature of the edge costs: the most costly edge could be any of the N_2 edges among the total number $N_1 + N_2$.

The result (3.60) generalizes to any graph $\mathcal{G}$, with edges $E(\mathcal{G}) = \{\epsilon_1, \ldots, \epsilon_n\}$ which have lengths of $N_1, \ldots, N_n$ and maximum costs $L_1, \ldots, L_n$ under a given lattice embedding. We fix an ordering $\pi'_{E(\mathcal{G})} = \pi'$, defined so that $L_{\pi'(1)} < L_{\pi'(2)} < \cdots < L_{\pi'(n)}$. Eq. (3.59) generalizes to

$$P_G(\ell_1, \ell_2, \ldots, \ell_n) = \Pr\left[\bigwedge_{i=1}^{n} (L_i \leq \ell_i)\right] = \prod_{i=1}^{n} p_i^{N_i}. \tag{3.62}$$

The generalization of (3.60) is then

$$\Pr[\pi' \wedge (G \leq \ell_0)] = \int_{-\infty}^{\ell_0} d\ell_{\pi'(n)} \int_{-\infty}^{\ell_{\pi'(n)}} d\ell_{\pi'(n-1)} \cdots \int_{-\infty}^{\ell_{\pi'(2)}} d\ell_{\pi'(1)} \prod_{i=1}^{n} \frac{d}{d\ell_i} \Pr\Big[\bigwedge_{i=1}^{n} (L_i \leq \ell_i)\Big]. \quad (3.63)$$

In the discussion which follows, I find it useful to identify the integro-differential operator

$$\mathcal{O}_{\text{MSF}}(\pi', \ell_0) = \int\cdots\int_{-\infty \leq \ell_{\pi'(1)} \leq \cdots \leq \ell_{\pi'(n)} \leq \ell_0} \prod_{i=1}^{n} d\ell_i \frac{d}{d\ell_i}. \quad (3.64)$$

We have

$$\sum_{\pi' \in S_n} \mathcal{O}_{\text{MSF}}(\pi', \ell_0) f(\ell_1, \ldots, \ell_n) = f(\ell_0, \ldots, \ell_0), \quad (3.65)$$

for any nonsingular function f of n parameters, since in summing over all the simplicial domains of integration we simply recover the integral of a total derivative. Equation (3.63) becomes

$$\begin{aligned} \Pr[\pi' \wedge (G \leq \ell_0)] &= \mathcal{O}_{\text{MSF}}(\pi', \ell_0) P_G(\ell_1, \ell_2, \ldots, \ell_n) \\ &= \prod_{i=1}^{n} \frac{N_{\pi'(i)}}{\sum_{j=1}^{i} N_{\pi'(j)}} p_0^{N_{\pi'(i)}} = \Pr[G \leq \ell_0] \prod_{i=1}^{n} \frac{N_{\pi'(i)}}{\sum_{j=1}^{i} N_{\pi'(j)}}. \end{aligned} \quad (3.66)$$

Using this result in (3.57) gives the desired expansion in terms of topological grpahs,

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(\ell_0) = \sum_{\mathcal{G} \in \mathcal{G}_{\mathbf{x},\mathbf{y},\mathbf{z}}} \sum_{\pi' \in S_{|E(\mathcal{G})|}} d_{\text{MSF}}(\mathcal{G}|\pi') \mathcal{O}_{\text{MSF}}(\pi', \ell_0) \sum_{\lambda : \mathcal{G} \to G} \Pr\Big[\bigwedge_{\epsilon \in E(\mathcal{G})} (L_\epsilon \leq \ell_\epsilon)\Big]. \quad (3.67)$$

This is the central result of this section, which should be compared with the percolation result (3.32). I now describe how a continuum theory may be constructed from this expansion.

3.3 A continuum perturbation expansion for MSF paths

In this section I describe how to turn the topological graph expansion obtained in the previous section into a continuum theory. I will argue that this may be accomplished by replacing the paths

of edges $\lambda(\epsilon)$ by continuum random walks and neglecting the self-avoidance and excluded-volume constraints on the set of allowable embeddings. Given these modifications, the only remnant of the lattice structure will be a short-distance (ultraviolet) cutoff. We expect the resulting perturbation series to be asymptotic, rather than convergent, as is typical for Feynman diagram expansions.

The continuum theory is defined by performing a modification procedure on each Feynman diagram of the field theory corresponding to bond percolation, and hence is defined to all orders in perturbation theory. Very roughly, we may think of the above procedure as "disassembling" each term the perturbation expansion for a percolation connectedness function into its histories of formation under the Kruskal process. However, the resulting expansion does *not* arise from the path integral of any action functional, so we do not technically have a "field theory of MSF paths," although I will show in the following section that many of the standard techniques of field theory remain applicable.

In section 3.3.1 I argue that the excluded volume constraint implicit in the lattice expansions may be safely neglected in the region around $d = 6$. I then describe how the lattice expansions of the previous section may be converted into a Feynman diagram series in the continuum in section 3.3.2. I describe how this series incorporates my existing mean-field scaling results in section 3.3.3, and in section 3.3.4 I describe how mean-field theory breaks down via a Ginzburg criterion analysis, confirming the value $d_c = 6$ obtained previously in section 2.2.7. This sets the stage for the renormalization group calculations I begin to set up in section 3.4.

3.3.1 The excluded volume constraint

A lattice expansion in terms of topological graphs, such as (3.32), is very close to describing a continuum theory. The only remaining roadblock lies in the sum over embeddings λ, which carries an effective excluded volume constraint: edges of $\mathcal{G}$ must be mapped to self-avoiding chains of lattice edges, and all these chains must be edge-disjoint (and hence also vertex-disjoint): no edge on the lattice may be used more than once. This is a highly non-trivial constraint to deal with in low dimensions. If we drop this constraint, we would have a sum over "free embeddings" $\bar{\lambda}$, which map edges of $\mathcal{G}$ to random walks in the lattice, which are allowed to intersect. This is the starting

point for a continuum theory: as is well known, the generating function for unconstrained random walks may be thought of as the propagator of a free scalar field (*e.g.*, [100]).

For percolation, the excluded volume constraint may be avoided by further modifications to the expansions given above. This is Essam's "ρ expansion", given in [40, 94], which takes the form

$$C_{\mathbf{x},\mathbf{y}}(\ell_0) = \sum_{\mathcal{G}\in\mathcal{G}_{\mathbf{x},\mathbf{y}}} d^*(\mathcal{G}) \sum_{\bar{\lambda}:\mathcal{G}\to G} \Pr[G \le \ell_0], \tag{3.68}$$

The d^* weights are defined in a manner similar to the d weights, and in fact $d(G) = d^*(G)$ for all graphs having no vertices of degree greater than three.

There certainly seems to be no obstacle to extending the ρ expansion to the expansion for MSF paths derived in the following section, but I do not pursue this line of inquiry (which may be relevant for a mathematically rigorous reformulation of the results given here). Instead, in what follows, I assume we may drop the excluded volume constraint without difficulty or modification of my lattice expansion. This is because my ultimate aim is a renormalization group calculation organized as an asymptotic series in $\varepsilon = d_c - d$. The dimensionality $d_c = 6$ is sufficiently high that the differences between self-avoiding paths and random walks will be negligible — or, more precisely, correspond to RG-irrelevant operators. Diagrammatically, imposing the self-avoidance constraints would lead to new graphs involving vertices of degree 4 or higher (arising from the intersection of two paths in Essam's expansion), which are RG-irrelevant interactions. For the remaining marginally relevant cubic interactions, we have $d^*(G) = d(G)$ and no modifications need to be made.

3.3.2 Continuum expansion for MSF paths

Continuum field theory of percolation

I begin by considering percolation connectedness functions, defined via (3.32) (and its generalization to n-point functions). Recall that these are equivalent to the corresponding MSF connectedness functions in the absence of any path vertex insertions, since the connectivity properties of the bond percolation and Kruskal processes are identical at any stage of cluster growth.

Since I argued in the previous section that we may safely neglect the excluded volume constraint present in the sum over embeddings, for a given topological graph $\mathcal{G}$ we may take the chains of edges produced by the lattice embedding to be independent random walks. The sum over embeddings may then be done by associating the propagator of a scalar field with each edge of $\mathcal{G}$; it is most natural to do this in terms of the Fourier transform with respect to the positions $\mathbf{x}_i$ of the embedded vertices of $\mathcal{G}$. The Fourier representation of the Green's function for the probability a random walk connects two vertices has the form $G_0(\mathbf{q}, t_0) = 1/(\mathbf{q}^2 + t_0)$, the same as the propagator of a scalar field [100]. Here the parameter $t_0 \geq 0$ corresponds to the square of the mass of the scalar field and arises from formally weighting all random walks according to length.

In scalar field theory, this mass can be viewed as arising from a quadratic interaction, which is the only RG-relevant interaction at $d = 6$. Therefore, in the field theory of percolation we should think of t_0 as depending monotonically on the percolation parameter p; naïvely, t_0 decreases to zero as p increases to p_c (however, this statement will be modified by perturbative corrections). Criticality corresponds to t_0, since this means the length scale set by t_0 diverges. This argument is borne out by the conventional derivation of the field theory of percolation via the Hubbard-Stratonovich transformation [89, 90].

The above considerations let us associate a Feynman diagram with each topological graph via the substitution

$$\sum_{\bar{\lambda}:\mathcal{G}\to G} \Pr[G \leq \ell_0] \to \frac{1}{\mathcal{A}(\mathcal{G})} I_{\mathcal{G}}(\mathbf{x}_1, \ldots, \mathbf{x}_n; t_0). \tag{3.69}$$

Here the symmetry factor $\mathcal{A}(\mathcal{G})$ arises from counting the automorphisms of $\mathcal{G}$ in the usual way (see discussion following eq. (3.32)), and $I_{\mathcal{G}}(\mathbf{x}_1, \ldots, \mathbf{x}_n; t_0)$ is a standard Feynman integral from scalar field theory, constructed by associating the propagator $G_0(\mathbf{x} - \mathbf{y}, t_0)$ of a scalar field to each edge of $\mathcal{G}$, and a factor of g_0 to each cubic vertex. The positions of the n external (root) vertices are given by $\mathbf{x}_1, \ldots, \mathbf{x}_n$. In the momentum basis this is

$$I_{\mathcal{G}}(\{\mathbf{k}_{\text{ext}}\}; t_0) = g_0^{-n} \prod_{\epsilon \in E(\mathcal{G})} \left(\int \frac{d^d \mathbf{q}_\epsilon}{(2\pi)^d} \frac{1}{\mathbf{q}_\epsilon^2 + t_0} \right) \prod_{v \in V(\mathcal{G})} g_0 (2\pi)^d \delta^d \left((\mathbf{k}_{\text{ext}})_v - \sum_{\epsilon \in E(\mathcal{G})} \mathcal{N}_{\epsilon, v} \mathbf{q}_\epsilon \right). \tag{3.70}$$

Here $\mathcal{N}$ is the incidence matrix of $\mathcal{G}$ under an arbitrary orientation of the edges:

$$\begin{aligned} \mathcal{N}_{e,v} &= 1 \text{ if } v \text{ is the head of } e, \\ &= -1 \text{ if } v \text{ is the tail of } e, \text{ and} \\ &= 0 \text{ otherwise.} \end{aligned} \tag{3.71}$$

The external momenta $\{\mathbf{k}_{\text{ext},i}\} i = 1, \ldots n$ are the Fourier conjugates of the positions $\{\mathbf{x}_i\}$ of the graph's root vertices, and $(\mathbf{k}_{\text{ext}})_v$ is the net external momentum flowing into vertex v. The momentum integrals should be done over the first Brillouin zone of the lattice we began with, which I denote by $\int_{\text{BZ}} d^d\mathbf{q}$, but as we are interested in long-distance (scaling) properties we will approximate this region by a sphere of radius Λ, which we will (for now) take to be of order $1/a$, where a is the lattice spacing. At this stage, g_0 is strictly speaking of order one, but will be viewed as small in the perturbation expansion (which may therefore be organized as an expansion in g_0).

Since the percolation d-weights are topologically invariant (section 3.1.4), the continuum expansion for the percolation connectedness functions then simply becomes

$$C_{\{\mathbf{x}_i\}}(\ell_0) \to \sum_{\mathcal{G} \in \mathcal{G}_{\{\mathbf{x}_i\}}} \frac{d(\mathcal{G})}{\mathcal{A}(\mathcal{G})} I_{\mathcal{G}}(\{\mathbf{x}_i\}; t_0). \tag{3.72}$$

I restrict the sum to topological graphs with vertices of degree three (as discussed above), except for the marked points $\mathbf{x}_i$, which are of degree one (thus we have g_0 for each cubic vertex). Both of these simplifying assumptions can be justified because other contributions can be shown to be RG-irrelevant near six dimensions, using arguments developed in section 3.4.

In section 3.1.3 I showed that the d-weights $d(\mathcal{G})$ are simply the combinatoric factors encountered in the standard treatment of percolation via the Potts model, and $\mathcal{G}_{\{\mathbf{x}_i\}}$ is the set of all cubic graphs with the given number of external legs, so this shows that this sector of my theory reproduces the standard field theory of bond percolation. The reader may compare my treatment with the conventional derivation via the Hubbard-Stratonovich transformation [89, 90]. I phrased my derivation in non-standard terms (not making use of an action functional) due to the fact that the MSF path vertex cannot be expressed in terms of any local operator: instead we must phrase the

argument to follow entirely in terms of diagrammatic expansions.

Mass insertion diagrams in percolation

We may define an analogous expansion for the function $C^e_{\mathbf{x},\mathbf{y}}(p)$ associated with singly-connected edges (defined in section 3.1.4). Here we sum over diagrams with two marked degree one external points at $\mathbf{x}$, $\mathbf{y}$, and with a single degree two vertex $\mathbf{z}$ marking the position of the single lattice edge e, as discussed above. This defines a set $\mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$. The appropriate symmetry factor is $\mathcal{A}'(\mathcal{G})$, and we pick up $|V(\mathcal{G})|-3$ factors of g_0 from the internal vertices. In the same way that this function on the lattice was obtained by differentiating with respect to the cost p_e of one edge, this function in the continuum can be obtained by differentiating the Feynman diagram expression for $C_{\mathbf{x},\mathbf{y}}(p)$ with respect to the mass t_ϵ associated with the edge $\mathbf{z}$ lies on (and no subsequent integration in the present case). In particular, this produces the correct symmetry factors. This procedure actually leads to the volume integral over $\mathbf{z}$ of the desired quality (*i.e.*, zero momentum flows through $\mathbf{z}$), but can be modified to allow some momentum to enter at $\mathbf{z}$, as given above. In field theory this operation is referred to as insertion of a mass- or ϕ^2 operator, where ϕ would be the field corresponding to the degree one external points [96].

Continuum theory of path vertex functions

I now consider the path connectedness function $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(p)$ for MSF(p), which (again) is the probability that at parameter value p there is a path on MSF(p) from $\mathbf{x}$ to $\mathbf{y}$ passing through $\mathbf{z}$. I treat the continuum version of the expansion in exactly the same way as I did for percolation, with the function $C^e_{\mathbf{x},\mathbf{y}}(p)$ discussed above being the closest analogue.

In the expansion for the path vertex connectedness functions (3.67) we needed to sum orderings π' of the costs L_ϵ of the edges ϵ of the diagram, weighted by appropriate factors. The first step was introducing the generalization

$$\Pr[G \le \ell_0] = \Pr\Big[\bigwedge_{\epsilon\in E(\mathcal{G})} (L_\epsilon \le \ell_0)\Big] \to \Pr\Big[\bigwedge_{\epsilon\in E(\mathcal{G})} (L_\epsilon \le \ell_\epsilon)\Big]. \tag{3.73}$$

I have argued above that when we pass to the continuum expansion the analogue of the parameter p_0 is the mass squared t_0, so in the continuum expansion we should proceed by giving each propagator an independent mass parameter t_ϵ:

$$\prod_{\epsilon \in E(\mathcal{G})} \frac{1}{\mathbf{q}_\epsilon^2 + t_0} \to \prod_{\epsilon \in E(\mathcal{G})} \frac{1}{\mathbf{q}_\epsilon^2 + t_\epsilon}. \tag{3.74}$$

The operator $\mathcal{O}_{\text{MSF}}(\pi, \ell_0)$ may trivially be rewritten in terms of the $\{t_\epsilon\}$ variables: because every derivative is paired with an integral, the Jacobians involved in changing from the $\{L_\epsilon\}$ to the $\{t_\epsilon\}$ cancel, and hence we do not require more detailed information on how the ts depend on the edge costs. The only difference is that ordering π is *reversed*: small L_ϵ corresponds to large t_ϵ; in Kruskal's algorithm on the lattice we approach p_c from below while here we approach t_c from above. Therefore, (3.64) becomes

$$\mathcal{O}_{\text{MSF}}(\pi, t_0) = \int \cdots \int_{\infty > t_{\pi(1)} \geq t_{\pi(2)} \geq \cdots \geq t_{\pi(n)} \geq t_0} \prod_{\epsilon \in E(\mathcal{G})} dt_\epsilon \frac{d}{dt_\epsilon}. \tag{3.75}$$

and the expansion for the MSF path vertex functions in the continuum is

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(\ell_0) \to \sum_{\mathcal{G} \in \mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}} \sum_{\pi' \in S_{|E(\mathcal{G})|}} \frac{d_{\text{MSF}}(\mathcal{G}|\pi')}{\mathcal{A}'(\mathcal{G})} \mathcal{O}_{\text{MSF}}(\pi', t_0) I_{\mathcal{G}}(\mathbf{x}, \mathbf{y}, \mathbf{z}; \{t_\epsilon\}). \tag{3.76}$$

The set of graphs involved are the same as those used in the mass-insertion graph of bond percolation (defined above), so we may compute this by starting with the expansion for $C^e_{\mathbf{x},\mathbf{y}}(t_0)$ and making the substitution

$$d(\mathcal{G}) I_{\mathcal{G}}(\mathbf{x}, \mathbf{y}, \mathbf{z}; t_0) \mapsto \sum_{\pi' \in S_{|E(\mathcal{G})|}} d_{\text{MSF}}(\mathcal{G}|\pi') \mathcal{O}_{\text{MSF}}(\pi', t_0) I_{\mathcal{G}}(\mathbf{x}, \mathbf{y}, \mathbf{z}; \{t_\epsilon\}) \tag{3.77}$$

on a diagram-by-diagram basis.

Summary

To summarize, the Feynman diagram rules for the MSF path vertex connectedness function, as specified in (3.77), are as follows:

1. Begin with the set of diagrams $\mathcal{G}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$ and the associated Feynman integrals $I_{\mathcal{G}}(\mathbf{x},\mathbf{y},\mathbf{z};t_0)$ defined as the Fourier transform of (3.70).

2. Associate a unique mass $t_0 \leq t_\epsilon < \infty$ to each propagator.

3. For each ordering $\pi' \in S_{|E(\mathcal{G})|}$ of these mass parameters, act on the integrand with the operator $\mathcal{O}_{\text{MSF}}(\pi', t_0)$ defined in (3.75).

4. Integrate over the momenta left undetermined by the δ-functions appearing in (3.70), subject to the cutoff $|\mathbf{q}_\epsilon| < \Lambda$.

5. Multiply the result of each integral by the appropriate diagrammatic weight $d_{\text{MSF}}(\mathcal{G}|\pi')/\mathcal{A}'(\mathcal{G})$, with d_{MSF} defined in (3.49) and $\mathcal{A}'(\mathcal{G})$ defined in section 3.1.4.

6. The sum of this quantity over all mass parameter orderings is the appropriate contribution to the MSF path connectedness function.

In sections 3.4.1, 3.4.2 I will reformulate these rules so as to be more amenable to actual computations. Note that we must act with $\mathcal{O}_{\text{MSF}}$ before any momentum integrations are performed, since the latter may produce expressions that diverge at the upper limit of the integrations in $\mathcal{O}_{\text{MSF}}$. This situation could alternatively be remedied by cutting off the domain of integration in $\mathcal{O}_{\text{MSF}}$ to $\{t_\epsilon\} < \Lambda^2$, at the expense of complicating my renormalization group calculation below.

Colloquially, we may think of the above procedure as "disassembling" each term the perturbation expansion for a percolation connectedness function into its constituent orderings or, roughly, histories of formation under the Kruskal process. This disassembly is accomplished by $\mathcal{O}_{\text{MSF}}$; we then reweight each history appropriately via the substitution $d \to d_{\text{MSF}}$ and then sum up all the pieces. This is *not* a literal explanation of the above procedure due to the fact that the expansions have been derived via inclusion-exclusion (some terms arise as corrections to overcounting in other terms) but it is perhaps the most intuitive explanation for the manipulations above.

3.3.3 Scaling and fractal dimensions: mean-field theory revisited

The perturbation expansion that we have now obtained can be organized as a loop expansion: the lowest order contribution to $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$ is order $\mathcal{O}(g_0^0)$, and is simply the diagram that takes the form of a path from $\mathbf{x}$ to $\mathbf{z}$ to $\mathbf{y}$, which possesses no loops (cycles), while higher orders in g_0 contain additional loops, one for each factor of g_0^2. The lowest order result, in position space and at $t_0 = 0$, takes the form (in this and the following, all separations like $|\mathbf{x} - \mathbf{y}|$ are assumed large, $\gg a$)

$$\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}} \propto \frac{1}{|\mathbf{x} - \mathbf{z}|^{d-2}|\mathbf{z} - \mathbf{y}|^{d-2}}. \tag{3.78}$$

By contrast, the 2-point connectedness function at criticality, obtained as a single scalar propagator, is proportional in this order to $1/|\mathbf{x} - \mathbf{y}|^{d-2}$. [At zero-loop order, these results are the same for MSF(p_c) and for critical percolation.] Dividing the two gives the conditional probability that there is a path from $\mathbf{x}$ to $\mathbf{y}$ passing through $\mathbf{z}$, given that there is a path from $\mathbf{x}$ to $\mathbf{y}$. Integrating over $\mathbf{z}$ gives $\propto |\mathbf{x} - \mathbf{y}|^2$. This is viewed as proportional to the total number of steps on the (lattice) path, even through the events of the path passing through the various z are not disjoint. The exponent 2 indicates that the fractal dimension of the path is 2, which is the correct result for a random walk. Thus I have shown that

$$D_{\mathrm{p}} = 2 \tag{3.79}$$

at zero-loop order. This will be found to be correct for $d > 6$, and also for $d = 6$ up to logarithmic corrections. The same dimension is believed to hold for paths on critical percolation clusters for $d > 6$, by similar field-theoretic arguments. Geometrically, it is because on large scales these clusters are trees, with no loops [39], and hence are the same in the MSF(p_c) process.

In general, and specifically for $d < 6$ as I will show, the scaling exponents will be different. At p_c, quite generally $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$ will have the scaling behavior

$$\widetilde{C}^{b\mathbf{z}}_{b\mathbf{x},b\mathbf{y}} = b^{-(d-2+\eta)-(d-D_{\mathrm{p}})}\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}} \tag{3.80}$$

for any $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{z}$, and scalar b, while the 2-point connectedness behaves as

$$C_{b\mathbf{x},b\mathbf{y}} = b^{-(d-2+\eta)} C_{\mathbf{x},\mathbf{y}}. \tag{3.81}$$

Thus these two functions determine two exponents η and $D_{\rm p}$ for MSF(p_c), and η will be the same as for percolation, as I will explain shortly ($\eta = 0$ for $d > 6$). Then in the same way as at zero loops, we infer the fractal dimension $D_{\rm p}$ for the MSF path. From the geometric point of view, $d - D_{\rm p}$ is the co-dimension of the path.

The exponent η describes the decay with distance r of the probability that two points are connected by a critical percolation cluster, namely $\sim r^{-(d-2+\eta)}$. In a field-theoretic point of view, $(d-2+\eta)/2 = x_\phi$ is the dimension of the Potts field operator ϕ, while $d - D_{\rm p} = x_{\rm p}$ is the scaling dimension for the path-vertex "operator". x_ϕ is related to the fractal dimension $D_{\rm perc}$ of the critical percolation clusters as the codimension $x_\phi = d - D_{\rm perc}$, so $D_{\rm perc} = (d + 2 - \eta)/2$.

For the percolation function $C^e_{\mathbf{x},\mathbf{y}}$, the corresponding ϕ^2 operator at e has dimension $x_{\phi^2} = d - D_{\rm sc}$, and $D_{\rm sc}$ is the fractal dimension of the set of singly-connected edges on the path from $\mathbf{x}$ to $\mathbf{y}$ on the critical percolation cluster [39]. For $d < 6$ this set does not usually form a connected path. We see that the MSF path must include the singly connected edges, which leads to the inequality

$$D_{\rm p} \geq D_{\rm sc}. \tag{3.82}$$

Because the function $C^e_{\mathbf{x},\mathbf{y}}$ is connected via differentiation with the change in connectivity with p (or t_0), the scaling dimension x_{ϕ^2} of the ϕ^2 insertion controls the length scale produced by taking $p < p_c$; this length is the correlation length ξ, and we can define the exponent ν by $\xi \sim (p_c - p)^{-\nu}$ as $p \to p_c$. It follows that $D_{\rm sc} = \nu^{-1}$ [39]. This discussion shows how the fractal dimension $D_{\rm p}$, and others, can be extracted from the renormalized perturbation calculations.

3.3.4 Beyond mean-field theory: breakdown of perturbation theory for $d < 6$

In section 3.4, I will prove that the perturbation expansion for the MSF path connectedness functions can be treated in a similar manner as that for standard field theories. In the current section,

at the risk of getting slightly ahead of myself, I will take these properties as given and describe the mechanism by which the mean-field scaling behavior of the previous section is corrected below d_c, and how these corrections come about due to the breakdown of the naïve perturbation series outlined above, which necessitates the introduction of the RG formalism.

Proper vertices

A first step is to introduce one-particle irreducible (1PI) functions. A diagram is defined to be 1PI if it does not become disconnected when a single edge is removed. Now for the MSF path connectedness function $\widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$, the (dominant) diagrams that contribute have a single edge emerging from $\mathbf{x}$ and $\mathbf{y}$. For terms of order $\mathcal{O}(g_0^2)$ (as $g_0 \to 0$), the diagram possesses at least one loop (cycle). It can then be decomposed into a chain of one or more disjoint 1PI 2-point graphs, connected by single edges. The vertex labeled $\mathbf{z}$ is either inside one of the subdiagrams (subgraphs), or on one of the single edges. The 1PI subdiagrams not containing $\mathbf{z}$ will be called self-energy diagrams.

For the d_{MSF}-weight of such a diagram, it is easy to see that the weights associated with each 1PI subdiagram factor. This is because the MSF path must pass through each of the 1PI subdiagrams in turn. For the self-energy diagrams, all paths through the subdiagram (which must be considered when evaluating the d_{MSF}-weight) contribute a non-zero amount (all diagrams we consider are connected to $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$). Consequently, the factor in the d_{MSF}-weight for the subdiagram is independent of the ordering π'_E restricted to the subdiagram, and then the weight reduces to the same expression as in percolation. (This is not true, however, for the 1PI subdiagram that contains the path vertex at $\mathbf{z}$.) The application of the $\mathcal{O}_{\text{MSF}}$ operator and the sum over orderings π'_E can now be carried out using (3.65). Then the contribution of such a self-energy diagram is the same as in percolation. The self-energy diagrams can be formally summed to all orders in perturbation theory to yield the self-energy $\Sigma(\mathbf{q}, t_0)$, and then each of the two series of alternating $G_0(\mathbf{q}, t_0)$'s and $\Sigma(\mathbf{q}, t_0)$'s can be summed as a geometric series, giving the full Green's function $G(\mathbf{q}, t_0)$,

$$G(\mathbf{q}, t_0)^{-1} = G_0(\mathbf{q}, t_0)^{-1} - \Sigma(\mathbf{q}, t_0) \tag{3.83}$$

(Dyson's equation). I pause to point out that the 2-point connectedness function for MSF(p), in which we do not require the path on the MSF to pass through any particular point $\mathbf{z}$, is similarly shown in this diagrammatic point of view to be the same as in percolation, and is given by $G(\mathbf{q}, t_0)$. Consequently, the exponent η defined above for $p = p_c$ must be the same as the similarly-defined exponent in percolation.

The MSF path connectedness function, with zero momentum entering at the path vertex, can now be written formally as

$$\int d^d\mathbf{z}\, \widetilde{C}^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}(t_0) = \int \frac{d^d\mathbf{q}}{(2\pi)^d} e^{-i\mathbf{q}(\mathbf{x}-\mathbf{y})}\, G(\mathbf{q}, t_0)\Gamma^{(2,\mathrm{PV})}(\mathbf{q}, \mathbf{0}; t_0)G(\mathbf{q}, t_0), \tag{3.84}$$

where $\Gamma^{(2,\mathrm{PV})}(\mathbf{q}, 0; t_0)$, which I call the path vertex function, is the Fourier transform of the sum of all 1PI diagrams with two external points (connected to $\mathbf{x}$, $\mathbf{y}$), plus the path vertex at $\mathbf{z}$, which has here been assigned zero momentum. (The generalization to $\Gamma^{(2,\mathrm{PV})}(\mathbf{q}_1, \mathbf{q}_2; t_0)$ should be obvious.) Diagrams contributing to $\Gamma^{(2,\mathrm{PV})}$ are depicted in figure 3.1. Similarly, I also define, for the N-point connectedness functions $G^{(N)}(\mathbf{q}_1, \ldots, \mathbf{q}_N)$ without the path vertex ($N = 2, 3$; $G^{(2)} = G$), and the 2-point connectedness function with a mass (or ϕ^2) insertion, which are the Fourier transforms of the percolation functions $G^{(2)} = C_{\mathbf{x},\mathbf{y}}$, $G^{(2,1)} = C^{\mathbf{z}}_{\mathbf{x},\mathbf{y}}$:

$$G^{(N)}(\{\mathbf{q}_i\}; t_0) = \prod_{i=1}^{N} G(\mathbf{q}_i, t_0)\Gamma^{(N)}(\{\mathbf{q}_i\}; t_0); \tag{3.85}$$

$$G^{(N,1)}(\{\mathbf{q}_i\}; \mathbf{q}; t_0) = \prod_{i=1}^{N} G(\mathbf{q}_i, t_0)\Gamma^{(N,1)}(\{\mathbf{q}_i\}; \mathbf{q}; t_0). \tag{3.86}$$

In these functions, a δ-function that sets the total wavevector to zero has been removed, and $\{\mathbf{q}_i\}$ stands for the ordered set $\mathbf{q}_1, \ldots, \mathbf{q}_N$. This causes a minor difference in notation from that for the path vertex function above: in the functions G or Γ containing N or $N + 1$ wavevector arguments, one of the wavevectors could be eliminated, which is what was done in $\Gamma^{(2,\mathrm{PV})}$ above, and I occasionally do this for the others also without further comment. The functions $\Gamma^{(N)}$ and $\Gamma^{(N,1)}$ are called the 1PI vertex functions (of the types indicated). I identify $\Gamma^{(2)}(\mathbf{q}_1, -\mathbf{q}_1; t_0) = G(\mathbf{q}_1; t_0)^{-1}$.

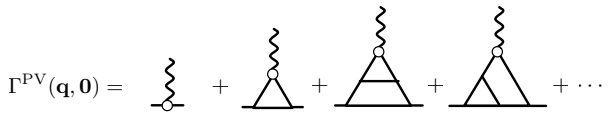

Figure 3.1: Perturbation expansion for the path vertex function given in (3.76). I represent the path vertex by the open circle connected to a wavy line. Note that these diagrams do not include any notation corresponding to the operation of $\mathcal{O}_{\mathrm{MSF}}$. The lowest non-trivial term is the second one on the right-hand side.

The problem of calculating the path exponent D_{p} has now been reduced to the calculation of the 1PI path vertex function $\Gamma^{(2,\mathrm{PV})}$. The external lines G are the same as in percolation, because of the factorization and ordering independence of the weights for the self-energy diagrams (which we are assuming for now). Similar, but more subtle, factorizations play an important role in the later part of the argument also. The path vertex, on the other hand, is not the same as the mass-insertion vertex in percolation which it resembles.

Ginzburg criterion and the breakdown of perturbation theory

To illustrate the perturbation expansion, let us now evaluate the first correction, of order g_0^2, to $\Gamma^{(2,\mathrm{PV})}$ (see Fig. 3.1) at zero external momentum. From the rules given above, this correction comes from the graph with three propagators connected to form a triangle, and the contribution is (note that $\mathcal{A}' = 1$ for this graph)

$$I_{\mathrm{MSF}}(\triangle, t_0) = \sum_{\pi' \in S_3} d_{\mathrm{MSF}}(\triangle|\pi)\mathcal{O}_{\mathrm{MSF}}(\pi, t_0) I_\Lambda(\triangle, \{t_1, t_2, t_3\}), \tag{3.87}$$

where the last integral is

$$I_\Lambda(\triangle, \{t_1, t_2, t_3\}) = g_0^2 \int^\Lambda \frac{d^d\mathbf{k}}{(2\pi)^d} \frac{1}{(\mathbf{k}^2 + t_1)(\mathbf{k}^2 + t_2)(\mathbf{k}^2 + t_3)}. \tag{3.88}$$

Here t_3 is the mass-squared on the edge not adjacent to the path vertex. Our momentum-space rules required us to integrate up to radius Λ. Since the diagram is evaluated at zero external momenta, it

is symmetric under permutations of the $\{t_\epsilon\}$ and

$$\mathcal{O}_{\rm MSF}(\pi, t_0) I_\Lambda(\triangle, \{t_1, t_2, t_3\}) = \frac{1}{3!} I_\Lambda(\triangle, \{t_0, t_0, t_0\}). \tag{3.89}$$

Applying (3.54), we have $d_{\rm MSF}(\triangle|\pi) = 0$ for the two orderings in which t_1 and $t_2 > t_3$, and $d_{\rm MSF}(\triangle|\pi) = -1$ for the other four orderings. Thus the result is

$$I_{\rm MSF}(\triangle, t_0) = -\frac{2}{3} I_\Lambda(\triangle, \{t_0, t_0, t_0\}). \tag{3.90}$$

For the corresponding simple mass-insertion vertex in percolation, the result would be instead $-I_\Lambda(\triangle, \{t_0, t_0, t_0\})$.

Dropping g_0^2 and numerical factors, the contribution to the path vertex function behaves like the integral

$$I'_\Lambda = \int_0^\Lambda \frac{k^{d-1} dk}{(k^2 + t_0)^3}. \tag{3.91}$$

At present, we are interested in this for fixed Λ as $t_0 \to 0$, so as to reach $p = p_c$. We see that for $d > 6$, (3.91) behaves as Λ^{d-6} as $\Lambda \to \infty$, and is finite (for any Λ) as $t_0 \to 0$. But for $d < 6$, the reverse is the case: the integral converges as $\Lambda \to \infty$, but diverges (for any Λ) as $t_0^{(d-6)/2}$ as $t_0 \to 0$. In the borderline case $d = 6$ the integral diverges logarithmically at both ends. For non-zero external wavevectors, the dependence of the integral on Λ is the same in all cases. Note that similar statements apply for percolation; only the numerical prefactor is different.

There are similar results for diagrammatic contributions with more loops. Simply counting the number of propagators and integrations gives the "superficial degree of divergence", which for the path vertex function is always $\sim g_0^2 \Lambda^{d-6}$ or $g_0^2 t_0^{(d-6)/2}$ for $d > 6$ and $d < 6$ respectively, raised to the power of the number of loops (independent cycles) in the diagram, as in the one-loop example above. Note that this is the same as if the $\mathcal{O}_{\rm MSF}$ operator were absent, because $\mathcal{O}_{\rm MSF}$ leaves the overall degree (in $\mathbf{k}$, at $t_0 \simeq 0$) of the integrand unchanged. The consequence for perturbation theory at fixed Λ and $d > 6$ is simple: each term in the perturbation expansion of $\Gamma^{(2,\rm PV)}$ for MSF(p) ($p \simeq p_c$) is finite as $t_0 \to 0$. This is true for the self-energy diagrams on the external lines also, and the value of $\Sigma(\mathbf{q}, t_0)$ at $\mathbf{q} = 0$ determines an effective shift in the value of t_0 that

corresponds to $p = p_c$: p_c must correspond to the value of t_0 such that $t_0 - \Sigma(\mathbf{0}, t_0) = 0$, and there are also other finite changes in the normalization of G. (In this case one would wish to sum up self-energy insertions in the lines inside of $\Gamma^{(2,\mathrm{PV})}$ also. However I have not shown that these take the same form as on the external lines. This will be addressed below.) But the consequence is that in each order, $\Gamma^{(2,\mathrm{PV})}(\mathbf{q}, t_0) \sim \mathcal{O}(1)$ as $\mathbf{q} \to 0$ at $p = p_c$, while $G(\mathbf{q}, t_0)^{-1} \propto \mathbf{q}^2$. This in turn implies that there is no change in the exponents from their lowest order values, $\eta = 0$ and $D_{\mathrm{p}} = D_{\mathrm{sc}} = 2$. Note that here I disregard the possibility that the sum of an infinite number of finite terms might diverge, which might invalidate the conclusion.

For $d \leq 6$, this perturbative argument breaks down as the corrections become arbitrarily large as $t_0 \to 0$, in particular in the region $t_0 < g_0^{4/(6-d)}$ (the Ginzburg criterion). In order to handle this, the use of RG techniques becomes essential. These techniques effectively re-sum and redefine the expansion. There are several formulations of the RG. These may be divided into two classes. One class of particularly powerful techniques is the field-theorists' RG, in which the aim initially is to take $\Lambda \to \infty$ (or $a \to 0$) at fixed separations or momenta, in such a way that the limits of the correlation (or connectedness) functions exist, thus recovering a true continuum theory. This is called renormalization of the theory. Subsequently, the renormalized theory is used to set up the RG, and calculate exponents for $d \leq d_c$. The leading alternative is the Wilsonian RG, in which the cutoff is kept finite. The Wilsonian RG is more difficult to use for higher numbers of loops. Both approaches lead to equivalent results for physical quantities such as exponents for $d \leq d_c$, where $d_c = 6$ for percolation and MSF(p_c). In this chapter I will follow the approach of the field theorists, since we want to make statements about the behavior of the expansion for arbitrary numbers of loops.

For $d > 6$, we can see from above that the effective expansion parameter is $g_0^2 \Lambda^{d-6}$. As $\Lambda \to \infty$, it is then necessary to make $g_0 \to 0$ such that $g_0^2 \Lambda^{d-6}$ does not diverge. In fact the situation is even worse than this would suggest: there simply is no rational way to define the limit $\Lambda \to \infty$ so that the connectedness functions at fixed $\mathbf{q}$ and (for example) $p = p_c$ have finite limits, without introducing an infinite number of parameters. This is referred to as *non-renormalizability* of the perturbation expansion. But by keeping the cutoff Λ finite, and using the Wilsonian point of

view, we can see that the exponents in this region take their simple zero-loop values, as indicated above. Accordingly, I concentrate on $d \leq 6$ from here on.

3.4 Renormalizability of the MSF field theory

In this section I give the proofs which establish that my perturbation expansion for MSF paths is renormalizable. Recall that diagrams of this theory with no path vertex are identical to those of percolation theory and hence pose no problem, while we construct diagrams involving the MSF path vertex by the substitution (3.77):

$$d(\mathcal{G})I(\mathcal{G}) \mapsto \sum_{\pi \in S_{|E(\mathcal{G})|}} d_{\text{MSF}}(\mathcal{G}|\pi)\mathcal{O}_{\text{MSF}}(\pi, t_0)I(\mathcal{G}, \{t_\epsilon\}), \tag{3.92}$$

where the integrals $I(\mathcal{G})$ and $I(\mathcal{G}, \{t_\epsilon\})$ are identical Feynman integrals with only cubic interaction vertices, containing the factor g_0 for each such interaction, but in the latter integral the mass-squared t_0 is generalized to a distinct parameter t_ϵ for each edge ϵ of the graph $\mathcal{G}$. In this section, I will drop the prime from the orderings π' throughout; orderings $\pi_{E(\mathcal{G})}$ are nonetheless the induced orderings on the set of highest costs L_ϵ of the set $E(\mathcal{G})$ of edges $\epsilon \in E(\mathcal{G})$ of a topological graph (Feynman diagram) or subgraph $\mathcal{G}$. d_{MSF} and $\mathcal{O}_{\text{MSF}}$ were defined in (3.54), (3.75) respectively, and as defined both depend on the structure of the entire graph $\mathcal{G}$. In particular, it is not entirely evident from the definition (3.54) how $d_{\text{MSF}}(\mathcal{G}|\pi)$ could be computed from knowledge of its values on subgraphs of $\mathcal{G}$.

Because the results and arguments are highly technical, I now give a more in-depth summary of the results obtained in this section. I begin in section 3.4.1 by introducing common terminology and explaining the parametric formulation of Feynman integrals. In section 3.4.2 I obtain the effect of the $\mathcal{O}_{\text{MSF}}$ operator, which is very simple in the parametric formulation: it introduces a simple product factor F_π into the integrand, which depends on the choice of an ordering π for the costs on the graph.

In section 3.4.3, I prove that the superficially-divergent subintegrations (as the cutoff $\Lambda \to \infty$) associated to a connected subgraph $\mathcal{H}$ come only from a subset $S' \subset S_{|E(\mathcal{G})|}$ of all possible

orderings on the edges of $\mathcal{G}$. Specifically, if $\mathcal{H}$ is a three-point subgraph or a 2-point subgraph containing the path vertex, then $S' = \pi_{[0]}$, in which all costs in the subgraph are cheaper than all those outside. Similarly, if $\mathcal{H}$ is a two-point subgraph (i.e. a self-energy), the only superficial divergences are for orderings $S' = \pi_{[0]} \cup \pi_{[1]}$, in which at most one edge in $\mathcal{H}$ has cost higher than one or more outside $\mathcal{H}$. Moreover, for these orderings, with one class of exceptions the divergences in self-energy or cubic coupling (three-point) subgraphs are the same as those in the corresponding percolation diagrams, up to the d-weights. These results generalize easily to diagrams with several superficially-divergent subdiagrams, if these are pairwise either disjoint or one inside another.

Having identified the important orderings, I consider in section 3.4.4 the behavior of the d_{MSF} weights for these orderings. I show that the weights obey nice factorization properties for connected subdiagrams with two or three external points for orderings in class $\pi_{[0]}$, and also (in a different, more general form) for self-energy subdiagrams $\mathcal{H}$ with orderings in which one or more edges in $\mathcal{H}$ is more costly than at least one outside $\mathcal{H}$. The factorization has the form

$$d_{\text{MSF}}(\mathcal{G}|\pi_{E(G)}) = d(\mathcal{H}) d_{\text{MSF}}(\mathcal{G}/\mathcal{H}|\pi_{E(\mathcal{G}/\mathcal{H})}) \tag{3.93}$$

if the path vertex is not in $\mathcal{H}$ (note the appearance of a d-weight from percolation), and

$$d_{\text{MSF}}(\mathcal{G}|\pi_{E(\mathcal{G})}) = d_{\text{MSF}}(\mathcal{H}|\pi_{E(\mathcal{H})}) d_{\text{MSF}}(\mathcal{G}/\mathcal{H}|\pi_{E(\mathcal{G}/\mathcal{H})}) \tag{3.94}$$

if the path vertex is in $\mathcal{H}$. The precise definitions, in particular for $\pi_{E(\mathcal{G}/\mathcal{H})}$, will be given in section 3.4.4. Here and below I use notation $\mathcal{G}/\mathcal{H}$ to denote the diagram obtained by contracting the subgraph $\mathcal{H}$ to a single vertex (which may be of degree 2, producing a harmless extension of the class of diagrams to be considered).

In section 3.4.5 I come to the heart of the proof. I use a theorem of Bergère and Lam [106] to show that the Feynman integral for each diagram in my perturbation expansion can be rendered absolutely convergent by a procedure of subtracting all the superficially divergent parts of the integrand. Furthermore, utilizing the results of preceding sections, all the terms that have to be subtracted for divergent subdiagrams (including those containing the path vertex) are the same as

those for a corresponding full diagram, with the exception of one class of terms as mentioned above, which is dealt with in section 3.4.6. Subject to the latter result, this means that all divergences are dealt with by renormalizing parameters and the overall scale of the vertex functions, as in a renormalizable field theory.

Finally, in section 3.4.6 I prove that the class of exceptional subleading divergences in the self-energy subdiagrams cancel in the sum over all diagrams of a given order. This completes the proof of renormalizability to all orders in the perturbation expansion.

3.4.1 Definitions

I begin by considering an arbitrary Feynman integral associated with a diagram $\mathcal{G}$ appearing in the perturbative expansion of a correlation function in, for example, the field theory of percolation.

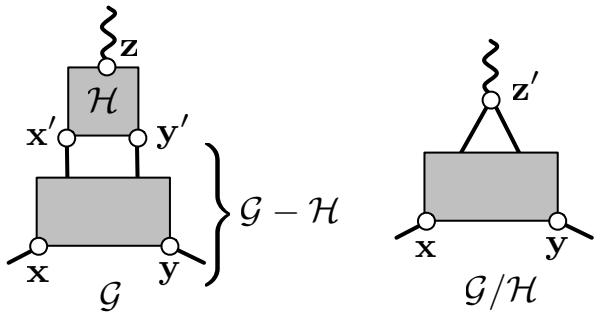

Figure 3.2: Depiction of an arbitrary graph contributing to the MSF path connectedness function. The various root vertices used in the definition of this function are labeled with open circles.

First, we recall the expression for a Feynman integral associated with $\mathcal{G}$ given in eq. (3.70). Let V, E be the vertex and edge sets of $\mathcal{G}$, and let $\mathcal{N}_{\epsilon,v}$ be its incidence matrix under an arbitrary orientation of its internal edges; i.e.

$$\begin{aligned} \mathcal{N}_{\epsilon,v} &= 1 \text{ if } v \text{ is the head of } \epsilon, \\ &= -1 \text{ if } v \text{ is the tail of } \epsilon, \text{ and} \\ &= 0 \text{ otherwise.} \end{aligned} \tag{3.95}$$

Neglecting the cutoff for a moment, the momentum integral with which we are concerned is

$$I_{\mathcal{G}}(\{\mathbf{k}_{\text{ext}}\}) = \prod_{\epsilon\in E(\mathcal{G})}\left(\int \frac{d^d\mathbf{q}_\epsilon}{(2\pi)^d}\frac{1}{\mathbf{q}_\epsilon^2+t_\epsilon}\right)\prod_{v\in V(\mathcal{G})}(2\pi)^d\delta^d\Big((\mathbf{k}_{\text{ext}})_v - \sum_{\epsilon\in E(\mathcal{G})}\mathcal{N}_{\epsilon,v}\mathbf{q}_\epsilon\Big). \tag{3.96}$$

Here $(\mathbf{k}_{\text{ext}})_v$ is the net external momentum incident on the vertex v. Since we will replace the percolation d-weight with the appropriate d_{MSF}-weight, we neglect the factor $d(\mathcal{G})$ and also g_0 to the power of the number of internal cubic vertices (or other couplings for interactions of different degree that may be present more generally).

Parametric representation of Feynman integrals

We make progress by expressing the Feynman integral (3.96) in terms of integrals over the Schwinger parameters α. This makes use of the identity $1/X = \int_0^\infty d\alpha\; e^{-\alpha X}$ to rewrite part of the integrand as

$$\prod_{\epsilon\in E}\frac{1}{\mathbf{q}_\epsilon^2+t_\epsilon} = \prod_{\epsilon\in E}\int_0^\infty d\alpha_\epsilon\; e^{-\alpha_\epsilon(\mathbf{q}_\epsilon^2+t_\epsilon)}. \tag{3.97}$$

(For brevity, I let $A_{\mathcal{G}}$ denote the set of parameters α_ϵ introduced above.) The total of $d|V|$ δ-functions can be rewritten using the identity $2\pi\,\delta(k) = \int d\lambda\, e^{i\lambda k}$ for each. Integrals over the internal momenta $\mathbf{q}_\epsilon$ are now Gaussian and can be performed, and then the λ-integrations become Gaussian and can be performed, except for one which produces a δ-function expressing conservation of the total momentum, $(2\pi)^d\delta^d(\sum_{v\in V}(\mathbf{k}_{\text{ext}})_v)$. Omitting this δ-function, we have

$$I_{\mathcal{G}}(\{\mathbf{k}_{\text{ext}}\}) = \frac{1}{(4\pi)^{\mathcal{L}d/2}}\int_0^\infty\prod_{\epsilon\in E}d\alpha_\epsilon\, P^{-d/2}(A_{\mathcal{G}})\exp\Big(-\mathbf{k}_{\text{ext}}^T\Delta^{-1}\mathbf{k}_{\text{ext}} - \sum_{\epsilon\in E}\alpha_\epsilon t_\epsilon\Big). \tag{3.98}$$

Here Δ is a $|V|\times|V|$ matrix, which is a Laplacian on $\mathcal{G}$, defined by

$$\Delta(A_{\mathcal{G}})_{v_1,v_2} = \sum_{\epsilon\in E}\mathcal{N}_{\epsilon,v_1}\frac{1}{\alpha_\epsilon}\mathcal{N}_{\epsilon,v_2}, \tag{3.99}$$

and $P(A_{\mathcal{G}})$ is defined as

$$P(A_{\mathcal{G}}) = \Big(\prod_{\epsilon\in E}\alpha_\epsilon\Big)\det{}'\Delta(A_{\mathcal{G}}), \tag{3.100}$$

in which the determinant $\det'$ is that of Δ with one row and column removed, so as to remove the zero mode. $\mathbf{k}_{\text{ext}}$ is viewed as a $|V|$-component vector, and $\mathcal{L} = \mathcal{L}(\mathcal{G})$ is the cyclomatic number of $\mathcal{G}$, the number of *independent* loops (cycles) of $\mathcal{G}$. $P(A_{\mathcal{G}})$ is a homogeneous polynomial of degree $\mathcal{L}$. These expressions are quite general and may be obtained for the diagrams of any field theory; for a further discussion consult [101, 107]. It is interesting that $P(A_{\mathcal{G}})$ and also Δ^{-1} can be related to weighted sums over spanning trees on $\mathcal{G}$ [101, 107] by the Kirchoff matrix-tree theorem [93]; it is not clear to us whether this fact is deeply involved in the renormalizability of the theory of MSF(p).

For many diagrams, the integral $I(\mathcal{G})$ as written in (3.96) or (3.98) is ultraviolet divergent and must be regularized, which is done by restricting the momentum integrations in (3.96) to the region $|\mathbf{q}_\epsilon| < \Lambda$. I implement this in (3.98) by taking the range of integration of each of the α_ϵ to be $[\Lambda^{-2}, \infty)$, which exponentially suppresses contributions from $|\mathbf{q}_\epsilon| \gg \Lambda$.

I make use of the parametric representation for Feynman integrals for two reasons. First, it greatly simplifies the study of renormalization of the expansion, as in the field theories in [101, 106, 108]. The second reason is that, as I show in the next section, the action of $\mathcal{O}_{\text{MSF}}$ takes a particularly simple form in this representation. Application of $\mathcal{O}_{\text{MSF}}$ directly to (3.96) results in intractable integrals over the $\{t_\epsilon\}$ for diagrams beyond one-loop order, while I am able to obtain its action on an arbitrary graph in closed form in equation (3.102).

3.4.2 Effect of the $\mathcal{O}_{\text{MSF}}$ operator

The preceding discussion applied to the diagrams from the field theory for percolation. To investigate how things change when we calculate MSF diagrams, we specify a given total ordering π of the masses of $\mathcal{G}$, such that

$$i < j \iff t_{\pi(i)} > t_{\pi(j)}. \tag{3.101}$$

The diagrammatic contribution to the MST theory is obtained by summing over all total orderings of edge costs consistent with the placement of the path vertex, according to (3.54). To find the contribution from one ordering π, we apply the operator $\mathcal{O}_{\text{MSF}}(\pi, t)$ defined in (3.75) to both sides

of (3.97), obtaining

$$\int_{\substack{\infty>t_{\pi(1)}>\cdots \\ \cdots>t_{\pi(|E|)}>t}} \cdots \int \prod_{\epsilon\in E} dt_\epsilon \, \frac{d}{dt_\epsilon} \frac{1}{\mathbf{q}_\epsilon^2 + t_\epsilon} = \int_{\Lambda^{-2}}^{\infty} \prod_{i=1}^{|E|} \frac{\alpha_{\pi(i)}\, d\alpha_{\pi(i)}}{\sum_{j=1}^{i} \alpha_{\pi(j)}} e^{-\sum_{\epsilon\in E} \alpha_\epsilon(\mathbf{q}_\epsilon^2 + t)}. \tag{3.102}$$

The integrand on the right-hand side is that appearing on the right-hand side of (3.97), multiplied by a factor

$$F_\pi(A_\mathcal{G}) \equiv \prod_{i=1}^{|E|} \frac{\alpha_{\pi(i)}}{\sum_{j=1}^{i} \alpha_{\pi(j)}}, \tag{3.103}$$

(I note that $F_\pi(A_\mathcal{G})$ is of the same form as I obtained on the lattice in equation (3.66). This is another manifestation of the well-known equivalence between scalar field theory and a system of random walkers.) Thus finally my prescription for evaluating the contribution of each diagram is that it is given by the parametric Feynman integral as for percolation, but with the factor

$$\sum_\pi d_{\text{MSF}}(\mathcal{G}|\pi) F_\pi(A_\mathcal{G}) \tag{3.104}$$

inserted inside the α integrals, replacing the $d(\mathcal{G})$ weight for the percolation theory.

The factor $F_\pi(A_\mathcal{G})$ obeys $0 \le F_\pi(A_\mathcal{G}) \le 1$ for any $A_\mathcal{G} \in [0,\infty)^{|E|}$, and has the property that it reduces to one as we go towards the limit in which

$$\alpha_{\pi(1)} \ll \alpha_{\pi(2)} \ll \cdots \ll \alpha_{\pi(|E|)}. \tag{3.105}$$

It tends to suppress orderings which do not obey the version of these inequalities in which all $\ll$'s are replaced by $<$'s. Thus it acts to replace the strict inequalities on the t_ϵ's by corresponding but softer conditions on the α_ϵ's. This result makes intuitive sense: high-momentum (small α) propagators correspond to lattice walks consisting of relatively few edges. In the Kruskal process, we expect the shortest paths to be completed first, at the lowest value of p, corresponding to a larger mass-squared t_0.

It will be useful to simplify the F_π factors as much as possible, by performing (or partially performing) the sums of $d_{\text{MSF}}(\mathcal{G}|\pi)F_\pi$ over orderings π as much as possible before performing the

integrals. I now give some basic formulas that are a step in this direction. First, I obtain another proof of (3.65) from the fact that

$$\sum_{\pi \in S_{|E|}} F_\pi(A_{\mathcal{G}}) = 1. \tag{3.106}$$

A more general fact that will be useful is that if we consider a subset of edges $E' \subseteq E$ of E and orderings π such that the masses on edges of E' are greater than all those in $E - E'$, and sum over all such orderings that fix an ordering on $E - E'$ (such orderings can be written as $\pi = \sigma \circ \pi_0$ for π_0 any one such ordering and σ a permutation in $S_{E'} \subset S_E$), then:

$$\sum_{\sigma \in S_{E'}} F_{\sigma \circ \pi_0}(A_{\mathcal{G}}) = \prod_{i=|E'|+1}^{|E|} \frac{\alpha_{\pi_0(i)}}{\sum_{j=1}^{i} \alpha_{\pi_0(j)}}, \tag{3.107}$$

in which the right-hand side is independent of the choice of π_0. This follows by using eq. (3.106) applied to the restricted sum over orderings. Indeed, as the derivation of this identity only used the sum over a smaller set, this can be used in a proof by induction (on the size $|E|$) of eq. (3.106) itself. The induction step, of taking $|E'| = |E| - 1$ and summing the right-hand side eq. (3.107) over cosets $S_{|E|}/S_{|E|-1}$ is simple.

3.4.3 Power counting for MSF Feynman integrals

In this section, I describe how the divergent behavior of a given diagram of the MSF path theory differs from that of the diagram from percolation theory from which it was obtained, and obtain some basic statements about the form of the divergences for each ordering.

In the absence of the F_π factor, the parametric form of the Feynman integrals may in general suffer from divergences associated with the region $\alpha \to 0$ for some or all α's. These take the place of the possibly more familiar divergences at large $\mathbf{k}$ in the original momentum space integrals over $\mathbf{q}_\epsilon$; recall that the latter integrals have already been done, after exchanging orders of integration. For a 1PI graph $\mathcal{G}$, the superficial degree of divergence of $I(\mathcal{G})$ is obtained easily from the momentum-space form by counting the total number of powers of all $\mathbf{q}_\epsilon$'s and integrations $\int d^d\mathbf{q}_\epsilon$, and is given by

$$\omega(\mathcal{G}) = d\mathcal{L}(\mathcal{G}) - 2|E(\mathcal{G})|. \tag{3.108}$$

This formula holds for any field theory of scalar fields interacting via non-derivative couplings. The same result is easily obtained in the parametric representation also [101, 106, 108]. It may be obtained more formally by rescaling the $\alpha_\epsilon \to \rho^2 \alpha_\epsilon$ for all edges of $\mathcal{G}$, with $\rho \to 0^+$. The formula may also be applied to the subintegral associated with a connected 1PI subdiagram $\mathcal{H}$ of $\mathcal{G}$ (strictly, a subdiagram is a subset of the vertices of $\mathcal{G}$, together with all edges that connect these vertices); this will be denoted $\omega(\mathcal{H})$. In this case, it is obtained from the behavior as the subset of α_ϵ associated with edges of $\mathcal{H}$ are scaled to zero by a common factor. Notice that the superficial degree of divergence for a subgraph $\mathcal{H}$ might be larger than that for $\mathcal{G}$. A graph or subgraph is said to be superficially divergent if its superficial degree of divergence is positive or zero, and superficially convergent if its superficial degree of divergence is negative. (A graph with $\omega = 0$ may diverge more slowly than any power of Λ, for example logarithmically, or may be convergent.) It is a theorem that if ω is negative for $\mathcal{G}$ and for all its subgraphs, then the associated Feynman (parametric) integral is absolutely convergent.

For the theory with cubic interactions that I consider here, the only connected 1PI graphs that are superficially divergent at $d = 6$ dimensions are (a) any self-energy diagram (with two external points), because all have $\omega = 2$, (b) any vertex correction diagram, that is a graph with three external points, because all have $\omega = 0$, and c) a self-energy graph with a ϕ^2 insertion, which have the same form as the vertex diagrams in b). Here in a) and b) an external point means that a line that "leaves" the graph (joined to it by a cubic vertex like the others) was removed to leave the 1PI part. The graphs containing ϕ^2 are relevant to the path vertex that we wish to consider in this chapter. Other graphs are superficially convergent.

Ordering dependence of the superficial degree of divergence

Turning to my theory for MSF(p), the parametric form of the Feynman integral for a given ordering π is simply modified by the insertion of the factor $F_\pi(A_\mathcal{G})$. Because $F_\pi(A_\mathcal{G})$ is bounded, it follows that the superficial divergence of any diagram or subdiagram of the MSF theory is no worse than the corresponding diagram of percolation theory from which it was obtained. More formally, F_π is a homogeneous rational function of degree zero, and so the superficial degree of

divergence for $\mathcal{G}$ is again $\omega(\mathcal{G})$.

However, for a subgraph $\mathcal{H}$ of $\mathcal{G}$, $F_\pi(A_\mathcal{G})$ may reduce the superficial degree of divergence below $\omega(\mathcal{H})$. Recall that for a subdiagram, we consider the limit as α_ϵ for $\epsilon \in E(\mathcal{H})$ go to zero simultaneously, by scaling them with a common factor ρ^2, leaving α_ϵ for $\epsilon \in \mathcal{G} - \mathcal{H}$ unchanged. Considering each of the $|E(\mathcal{H})|$ factors in $F_\pi(A_\mathcal{G})$ that have numerator α_ϵ for an edge $\epsilon \in E(\mathcal{H})$ appearing in (3.103) in this limit, we see that in this limit $\rho \to 0$,

$$F_\pi(A_\mathcal{G}) = \mathcal{O}\left(\rho^{2n_\pi(\mathcal{H},\mathcal{G})}\right), \tag{3.109}$$

where I define $n_\pi(\mathcal{H},\mathcal{G})$ to be the number of masses t_ϵ for $\epsilon \in E(\mathcal{H})$ that are less than at least one of the masses in $E(\mathcal{G})$ under the ordering π. (Clearly $n_\pi(\mathcal{H},\mathcal{G}) = 0$ for $\mathcal{H} = \mathcal{G}$.) For a fixed 1PI connected subgraph $\mathcal{H}$ of $\mathcal{G}$, this provides a useful partitioning of orderings into sets $\pi_{[m]}$, $m = 0$, $1, \ldots$:

$$\pi_{[m]} = \{\pi : n_\pi(\mathcal{H},\mathcal{G}) = m\}. \tag{3.110}$$

Thus the orderings $\pi_{[0]}$ (which will prove most important in what follows), for which $F_\pi = \mathcal{O}(1)$ as $\rho \to 0$, are those where all of the masses on the edges of the subgraph $\mathcal{H}$ are larger than those in $\mathcal{G} - \mathcal{H}$, that is all the costs in $\mathcal{H}$ are lower.

We may add this result to the superficial degree of divergence to obtain the overall superficial degree of divergence of a connected 1PI subdiagram $\mathcal{H}$ of a connected 1PI diagram $\mathcal{G}$ under the ordering π:

$$\omega_{\mathrm{MSF}}(\mathcal{H},\mathcal{G}|\pi) = \omega(\mathcal{H}) - 2n_\pi(\mathcal{H},\mathcal{G}). \tag{3.111}$$

This implies that it is only for class $\pi_{[0]}$ that the superficial degree of divergence of the subgraph $\mathcal{H}$ is unchanged by F_π. For subgraphs with $\omega(\mathcal{H}) = 0$ (i.e. the vertex and path vertex diagrams), orderings other than those in $\pi_{[0]}$ give convergent subintegrals. For the self-energy subgraphs, with $\omega = 2$, orderings in $\pi_{[1]}$ lower the superficial degree of divergence to 0, and these are additional divergences with which we will have to deal. Moreover, in all cases there are subleading terms in the behavior of F_π as $\rho \to 0$ for a subgraph, and while these terms are superficially convergent in most cases, the first subleading term also has zero superficial degree of divergence in the case of

the self-energy subdiagrams.

For further analysis, it is helpful to consider the sum $\sum_\pi d_{\text{MSF}}(\mathcal{G}|\pi)F_\pi(A_\mathcal{G})$ and to attempt to simplify it as much as possible, so that the evaluation of the parametric integrals reduces to those for percolation as much as possible. Indeed, by the "contribution of a diagram" in general I mean the weighted sum over orderings. In order to consider divergent subintegrals for subdiagrams, it is useful to have factorization properties of the weights d_{MSF}. I consider these properties next.

3.4.4 Factorization properties of MSF diagrammatic weights

In this section I demonstrate that the diagrammatic weights d_{MSF} possess enough factorization properties for my proof of the renormalizability of the perturbation expansion to go through. Let us first recall that for the d-weights in percolation, the weight for a diagram $\mathcal{G}$ containing a 2- or 3-point subdiagram $\mathcal{H}$ factors into the weight for $\mathcal{H}$ times that for the "quotient graph" $\mathcal{G}/\mathcal{H}$ in which the subgraph $\mathcal{H}$ is shrunk to a single vertex (formally, its vertices are identified, and its edges are deleted): $d(\mathcal{G}) = d(\mathcal{H})d(\mathcal{G}/\mathcal{H})$. This is immediate in the Potts model formulation in which the d-weights originate from contracting together tensors, due to S_Q permutation symmetry (apart from the problem of giving a formal definition of the $Q \to 1$ limit). It can also be derived from the combinatorial definitions described in section 3.1.3 (this is shown in the case of some 2-point subdiagrams in Ref. [92]). It is important for the proof of renormalizability, as the contributions of such subgraphs in Feynman integrals will be treated as "correcting" or "renormalizing" the parameters attached to 2- and 3-point vertices in the graphical expansion. We require some similar properties in the expansion for MSF(p_c).

1PI decomposition

Recall that the weights can be defined as in eq. (3.54) [for $E = E(\mathcal{G})$],

$$d_{\text{MSF}}(\mathcal{G}|\pi) = \sum_{E' \subseteq E} (-1)^{|E|-|E'|} \mathbb{I}[E' \text{ connects } \mathbf{x}, \mathbf{y}] \mathbb{I}_c^{(\mathbf{z})}[\gamma_{\text{MST}}(\mathcal{G}_{E'}|\pi_{E'})]. \tag{3.112}$$

This differs from the diagrammatic weight for percolation (3.13) only in the presence of the additional indicator function $\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\text{MST}}(\mathcal{G}_{E'}|\pi_{E'})]$. A graph that is not 1PI can be decomposed into (connected) 1PI subdiagrams lying on a chain of single edges and such 1PI parts that form a path from $\mathbf{x}$ to $\mathbf{y}$, and possibly other 1PI parts. That is, $\mathcal{G}$ may be constructed as a tree $\mathcal{G}_0$ (with $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$ marked) which is then decorated by replacing its vertices v with subgraphs $\mathcal{H}_v$. As the MSF path must pass through a chain of 1PI parts, it follows that for those 1PI subdiagrams that do not contain the vertex at $\mathbf{z}$, the indicator $\mathbb{I}_c^{(\mathbf{z})}$ is independent of the path through such a 1PI subdiagram, and accordingly the d_{MSF} weight factors into a product of weights for the single edges and for the 1PI parts. Moreover the d_{MSF} factor for each such 1PI subdiagram reduces to d in percolation for that subdiagram (for a single edge, the weight is 1). Likewise, for a vertex-reducible subdiagram or "tadpole", such as a 1PI part connected to the rest by a single edge, the d_{MSF}-weight is the same as in percolation and vanishes. Similarly the d_{MSF} weight for a diagram that contains a disconnected subdiagram vanishes. Hence from here on we need consider only connected, vertex-irreducible 1PI diagrams $\mathcal{G}$ that contain the path vertex at $\mathbf{z}$, as well as root points that we can relabel as $\mathbf{x}$, $\mathbf{y}$.

Ordering dependence

For MSF(p), the weights d_{MSF} depend on the ordering π of the costs of the edges of the topological graph $\mathcal{G}$, as well as on $\mathcal{G}$. In this section, I will denote these costs by the original symbol ℓ_ϵ for edge $\epsilon \in \mathcal{G}$ (these costs in fact stand for the maximum, earlier denoted L_ϵ, of the chain of edges that are the image of ϵ under an embedding of $\mathcal{G}$ in the lattice). In terms of the costs, the ordering π is defined by

$$i < j \iff \ell_{\pi(i)} < \ell_{\pi(j)}. \tag{3.113}$$

(I use the costs, rather than the mass-squared's for which the inequalities are reversed, because I find the former arrangement to be more intuitive.) In seeking a factorization similar to that for the d-weights in percolation, there are two issues. Because the d_{MSF}-weights depend on a choice of ordering, one issue is whether some factorization holds at all for each ordering, and a second is, if there is some factorization, what ordering would be used for the quotient $\mathcal{G}/\mathcal{H}$. What I obtain below may not be the most general possible result. Instead I obtain statements for two (overlapping) sets

of conditions, and these are sufficient for my purposes.

Factorization for orderings in $\pi_{[0]}$

Motivated by the considerations of which orderings produce ultraviolet-divergent Feynman integrals associated with a subgraph, I first show that for orderings π such that all edges in a connected subgraph $\mathcal{H}$ have lower cost than all others in $\mathcal{G}$, where $\mathcal{H}$ is a 2- or 3-point subgraph, and in the 2-point case the vertex $\mathbf{z}$ can also be present, factorization holds:

$$d_{\text{MSF}}(\mathcal{G}|\pi) = d_{\text{MSF}}(\mathcal{H}|\pi_{\mathcal{H}}) d_{\text{MSF}}(\mathcal{G}/\mathcal{H}|\pi_{\mathcal{G}-\mathcal{H}}). \tag{3.114}$$

Here the right-hand side involves the ordering $\pi_{\mathcal{H}}$, which is π restricted to $\mathcal{H}$, and $\pi_{\mathcal{G}-\mathcal{H}}$ which is π restricted to $\mathcal{G} - \mathcal{H}$. (For graphs $\mathcal{G}$, subgraphs $\mathcal{H}$, and quotients $\mathcal{G}/\mathcal{H}$, I will allow abuses of notation like $\pi_{E(\mathcal{H})} = \pi_{\mathcal{H}}$.) Further, in the case in which $\mathcal{H}$ does not contain $\mathbf{z}$, we already know that $d_{\text{MSF}}(\mathcal{H}|\pi_{\mathcal{H}}) = d(\mathcal{H})$. We recall that these orderings are those in class $\pi_{[0]}$, which produce the *leading* divergence for the 2- or 3-point subgraphs. Further, the factorization generalizes to the case in which there are several disjoint such subgraphs, and the costs in the union of the sets of edges of the subgraphs are lower than those in the remainder of $\mathcal{G}$ (regardless of the relative orderings among the edges in the subgraphs). In this case, each disjoint subgraph carries a weight as for the single subgraph considered above. Then, because the 2-point (or self-energy) subgraph (that does not contain $\mathbf{z}$) also has subleading divergences that occur when its costs do not obey the preceding conditions, I also derive a more general result for such a subgraph for any ordering. These results can be combined to handle a large class of orderings and subgraph structures.

First I show that, if $\mathcal{H}$ is a 2- or 3-point subgraph, then in the sum over subsets of edges $E' \subseteq E$ in d_{MSF} we can replace

$$\mathbb{I}[E' \text{ connects } \mathbf{x}, \mathbf{y}] = \mathbb{I}[E'(\mathcal{H}) \text{ connects } \{\mathbf{x}_i\}] \mathbb{I}[E'(\mathcal{G}/\mathcal{H}) \text{ connects } \mathbf{x}, \mathbf{y}], \tag{3.115}$$

where $\mathbf{x}, \mathbf{y}$ are the root points of $\mathcal{G}$, and $\{\mathbf{x}_i\}$ are the root points of $\mathcal{H}$, because other terms cancel. To see this, first notice that if for a 2- or 3-point subgraph $\mathcal{H}$, the "diluted" edge set of $\mathcal{H}$, $E'(\mathcal{H}) =$

$E' \cap E(\mathcal{H})$ [and similarly for $E'(\mathcal{G}/\mathcal{H})$], does not connect all the root vertices, then there is at least one root vertex not connected to any of them (this does not hold for a subgraph with more than three root points). Choose one of these, and without loss of generality suppose it is $\mathbf{x}_1$. In $\mathcal{G}$ there is a single edge incident on $\mathbf{x}_1$ that is not in $\mathcal{H}$ (call it $\epsilon_{\mathbf{x}_1}$). The minimum spanning tree path $\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})$ from $\mathbf{x}$ to $\mathbf{y}$ on E' clearly cannot pass though $\mathbf{x}_1$ for such an E', whatever the ordering π. We can pair off such subsets E' by choosing pairs of E' which are the same subsets except that the edge $\epsilon_{\mathbf{x}_1}$ is in one and not in the other. These subsets differ in size by one, and the indicator function $\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})]$ takes the same value for both. Hence these contributions cancel, and the result follows.

Now I turn to the factoring of $\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})]$; it is here that the form of the ordering enters. The case in which all edges in the subgraph $\mathcal{H}$ have costs lower than all those in $\mathcal{G}-\mathcal{H}$ is quite simple. First, the same property is inherited in the ordering $\pi_{E'}$ restricted to E'. As the Kruskal process runs on E', these edges are tested first, and when that is completed the root points of the subgraph are connected (this follows because I have shown that E' connects these vertices). For the remainder of the process, from which the path $\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})$ is obtained, the subgraph $\mathcal{H}$ [or its diluted version which I denote $\mathcal{H}_{E'(\mathcal{H})}$] can be viewed as collapsed to a single vertex to produce $\mathcal{G}_{E'}/\mathcal{H}_{E'(\mathcal{H})}$. It is useful now to distinguish two cases: either $\mathbf{z}$ is in $\mathcal{H}$, or it is not. In the first case, the MSF path must enter $\mathcal{H}$ to reach $\mathbf{z}$, and then leave. This implies that a) on $\mathcal{G}_{E'}/\mathcal{H}_{E'(\mathcal{H})}$, the image of $\mathcal{H}$ is the point through which the MSF path must pass, and b) once within $\mathcal{H}$ the path must pass through $\mathbf{z}$. That is, we can write for the indicator function

$$\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})] = \mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'(\mathcal{H})})] \times \mathbb{I}_c^{(\mathcal{H})}[\gamma_{\rm MST}(\mathcal{G}_{E'}/\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'-E'(\mathcal{H})})] \quad (3.116)$$

The summation over subsets of the edges E' can be written as a sum over subsets $E'(\mathcal{H})$ and over $E'' = E' - E'(\mathcal{H})$, and so the factorization of the $d_{\rm MSF}$-weights as in eq. (3.114) follows. Likewise, in the case where the path vertex is located in $\mathcal{G}-\mathcal{H}$, we can simply write

$$\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})] = \mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}/\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'-E'(H)})], \quad (3.117)$$

and again the form in eq. (3.114) follows, though now $d_{\rm MSF}(\mathcal{H}|\pi_{\mathcal{H}}) = d(\mathcal{H})$. Together these prove all the relations shown in Fig. 3.3 for the the stated class of orderings.

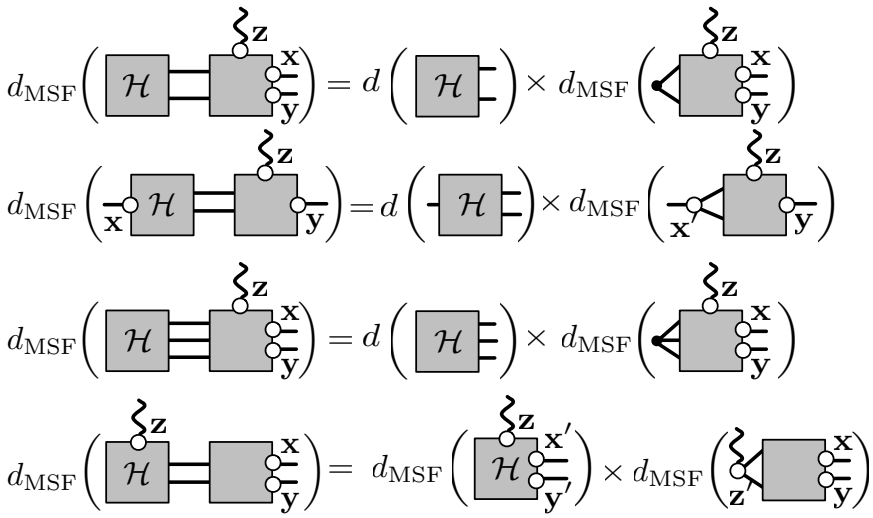

Figure 3.3: Factorization properties of $d_{\rm MSF}(\mathcal{G}|\pi)$ for connected subgraphs that are proved in the text for orderings in the class $\pi_{[0]}$.

For graphs with several disjoint subgraphs of the same type as the single connected subgraphs considered above, the same proof goes through if all edges in all the subgraphs are less costly than those outside. That is, the edges in one of the connected subgraphs need not be all more, nor all less, costly than those in another one of the subgraphs. Thus this result is more general than simply iterating the application of the preceding result, though the final factored form of the $d_{\rm MSF}$-weight is the same as if it were.

Arbitrary ordering: two-point subgraph case

I now turn to a more general argument for the case of a connected 1PI 2-point subgraph $\mathcal{H}$ and *any* ordering π. It holds if the path vertex $\mathbf{z}$ is within $\mathcal{H}$, however for orderings not in the class $\pi_{[0]}$ which are already covered by the preceding proof, the corresponding Feynman integrals are convergent, so I will not make of this, and can assume that $\mathbf{z}$ is not in $\mathcal{H}$.

Let $\epsilon_{\mathcal{H}}$ be the most costly edge in $\mathcal{H}$ for the ordering π. Now we consider the evaluation of the $d_{\rm MSF}$-weight. For each diluted edge-set E', this involves comparison of paths from $\mathbf{x}$ to $\mathbf{y}$ (which exist because the contribution to $d_{\rm MSF}$ vanishes if $\mathbf{x}$ and $\mathbf{y}$ are not connected). We saw above that

we can assume that the root points $\mathbf{x}_1$, $\mathbf{x}_2$ of $\mathcal{H}$ are connected by $E'(\mathcal{H})$. Further, there are paths from $\mathbf{x}$ to $\mathbf{y}$ through $\mathcal{H}_{E'(\mathcal{H})}$, because otherwise $\mathcal{H}_{E'(\mathcal{H})}$ is either disconnected from both $\mathbf{x}$ and $\mathbf{y}$, or is part of a tadpole, and in either case the weight vanishes as we saw above. To find the MST path from $\mathbf{x}$ to $\mathbf{y}$, the task can be broken into subtasks, and one of these is first to find the MST path through $\mathcal{H}_{E'(\mathcal{H})}$ between its root points. If the MST path from $\mathbf{x}$ to $\mathbf{y}$ passes through $\mathcal{H}_{E'(\mathcal{H})}$, the portion within $\mathcal{H}_{E'(\mathcal{H})}$ must be this MST path. I now show that (within the sum defining d_{MSF}) this path $\gamma_{\text{MST}}(\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'(\mathcal{H})})$ between $\mathbf{x}_1$ and $\mathbf{x}_2$ must pass through the most costly edge $\epsilon_{\mathcal{H}}$ of $\mathcal{H}$. For suppose that $\epsilon_{\mathcal{H}} \in E'(\mathcal{H})$, but the MST path does not traverse it. Then there is another edge set which is the same as E' except that $\epsilon_{\mathcal{H}}$ is omitted, and these terms cancel in pairs (note that the MST paths are the same for these edge sets). But the terms with $\epsilon_{\mathcal{H}} \in E'(\mathcal{H})$ and $\epsilon_{\mathcal{H}}$ on $\gamma_{\text{MST}}(\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'(\mathcal{H})})$ do not cancel in a similar way, as removing $\epsilon_{\mathcal{H}}$ from this edge set leaves the root vertices $\mathbf{x}_1$ and $\mathbf{x}_2$ disconnected, and we know that those edge sets cancel among themselves. The reason the root vertices become disconnected on removing $\epsilon_{\mathcal{H}}$ (so $\mathcal{H}_{E'(\mathcal{H})}$ is not 1PI) is that if not, then a less-costly path (in the sense of the ordering $\prec$ in section 3.2.2) between the roots would exist.

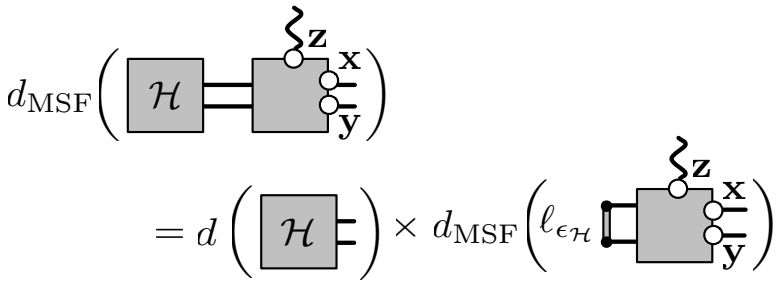

Figure 3.4: The factorization property of $d_{\text{MSF}}(\mathcal{G}|\pi)$ proved in the text for arbitrary ordering of edge costs.

It follows that in comparing possible MST paths on $\mathcal{G}_{E'}$, the subgraph $\mathcal{H}_{E'(\mathcal{H})}$ can be replaced by a single edge from $\mathbf{x}_1$ to $\mathbf{x}_2$ with cost $\ell_{\epsilon_{\mathcal{H}}}$. I use this result to define the induced ordering $\pi_{\mathcal{G}/\mathcal{H}}$ for the quotient graph for such a 2-point subgraph $\mathcal{H}$; this ordering gives the ordering for any diluted edge set $E'(\mathcal{G}/\mathcal{H})$. Note however that here we are forced to view $\mathcal{H}$ as replaced by an edge, not a vertex, in the quotient graph. [Further, $\mathcal{H}$ is bordered by two other edges, and these three edges form a chain, which by the general elementary arguments given earlier can be replaced

by a single edge of cost the maximum of the costs of the three edges, for the purposes of finding the MST path $\gamma_{\rm MST}(\mathcal{H}_{E'}|\pi_{E'})$ from $\mathbf{x}$ to $\mathbf{y}$.] We can summarize this whole argument as showing that the indicator function can be written as

$$\mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}|\pi_{E'})] = \mathbb{I}_c^{(\epsilon_{\mathcal{H}})}[\gamma_{\rm MST}(\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'(\mathcal{H})})] \times \mathbb{I}_c^{(\mathbf{z})}[\gamma_{\rm MST}(\mathcal{G}_{E'}/\mathcal{H}_{E'(\mathcal{H})}|\pi_{E'(\mathcal{G}/\mathcal{H})})], \quad (3.118)$$

while the edge subsets $E'(\mathcal{G}/\mathcal{H})$ that have to be summed over are subsets of the set $(E' - E'(\mathcal{H})) \cup \{\epsilon_{\mathcal{H}}\}$ (the latter change cause no difficulty, and again the three edges in a chain can be replaced by one, with the cost as described above). This then shows that the weight factors as

$$d_{\rm MSF}(\mathcal{G}|\pi) = d_{\rm MSF}(\mathcal{H}|\pi_{\mathcal{H}}) d_{\rm MSF}(\mathcal{G}/\mathcal{H}|\pi_{\mathcal{G}/\mathcal{H}}), \quad (3.119)$$

where the $d_{\rm MSF}$ for $\mathcal{H}$ is that for the path from $\mathbf{x}_1$ to $\mathbf{x}_2$ to pass though $\epsilon_{\mathcal{H}}$, while the second simply requires a path on the quotient graph $\mathcal{G}/\mathcal{H}$ to pass through $\mathbf{z}$. However, the argument already given above for the MST path within $\mathcal{H}$ on the diluted edge sets $E'(\mathcal{H})$ shows that if the former condition is dropped, then the evaluation of the sum is the same. That is

$$d_{\rm MSF}(\mathcal{H}|\pi_{\mathcal{H}}) = d(\mathcal{H}) \quad (3.120)$$

for $d_{\rm MSF}(\mathcal{H}|\pi_{\mathcal{H}})$ with the MST path vertex at the most costly edge of $\mathcal{H}$ under the ordering $\pi_{\mathcal{H}}$. Thus we obtain factorization in the same form as before, as desired (see Fig. 3.4). This agrees with the result for an ordering in class $\pi_{[0]}$ (because then $\pi_{\mathcal{G}/\mathcal{H}} = \pi_{\mathcal{G}-\mathcal{H}}$), but gives the correct generalization to other orderings, for the case of a 2-point subgraph. For other orderings, the highest cost in $\mathcal{H}$ has to be compared with those in the remainder of $\mathcal{G}$; I emphasize again this aspect of the definition of $\pi_{\mathcal{G}/\mathcal{H}}$.

In the present case, the argument can simply be used again if $\mathcal{G}/\mathcal{H}$ contains a 2-point subgraph. For 3-point subgraphs, I expect that a more complicated generalization exists, but I have not looked for it.

3.4.5 Proof of renormalizability

In this section, I assemble the preceding results to describe the divergences of the diagrams or subdiagrams in MSF(p) theory, and compare them with the corresponding ones in the perturbation expansion for percolation. Here by a diagram, I mean the corresponding Feynman integral, including the sum over orderings of the $d_{\rm MSF}F_\pi$ factors inside the parametric integral. The initial results provide the direct motivation for the renormalization of the perturbation series. Then I describe the proof of renormalizability.

Summary of results so far

We saw in section 3.4.3 that for a subdiagram $\mathcal{H}$ and an ordering in the class $\pi_{[0]}$ (or for the whole diagram $\mathcal{G}$, and any ordering), the superficial degree of divergence $\omega_{\rm MSF}(\mathcal{H},\mathcal{G}|\pi)$ is the same as $\omega(\mathcal{H})$. I consider only 2- or 3-point subdiagrams, including the 2-point subdiagram that contains the path vertex. For any fixed ordering on the edges $E(\mathcal{G})-E(\mathcal{H})$ not in $\mathcal{H}$, we can consider the sum of $d_{\rm MSF}(\mathcal{G}|\pi)F_\pi(A_\mathcal{G})$ over all the orderings $\pi_\mathcal{H}$ of edges in $E(\mathcal{H})$ such that the ordering of all edges is in $\pi_{[0]}$. For each ordering in the sum, we saw in section 3.4.4 that the $d_{\rm MSF}$ weight reduces to $d_{\rm MSF}(\mathcal{H}|\pi_\mathcal{H})d_{\rm MSF}(\mathcal{G}/\mathcal{H}|\pi_{\mathcal{G}-\mathcal{H}})$ (for the path vertex case) or $d(\mathcal{H})d_{\rm MSF}(\mathcal{G}/\mathcal{H}|\pi_{\mathcal{G}-\mathcal{H}})$ (for the other cases). For the latter cases in which the path vertex is not in $\mathcal{H}$, the weight $d_{\rm MSF}(\mathcal{G}|\pi)$ is independent of the ordering $\pi_\mathcal{H}$, and the sum over the latter can be performed using eq. (3.107), which shows that the part of the F_π associated with the subdiagram has reduced to unity, as in eq. (3.106). The remaining factor on the right hand side of eq. (3.107) depends on the α's for the subdiagram, but of course not on their ordering. To leading order as all those α's are scaled to zero, the resulting subintegration has exactly the same divergence (not only degree of divergence) as the corresponding subdiagram in the field theory of percolation near criticality [including the $d(\mathcal{H})$ weight], and the remaining F_π factor is that for the quotient graph, $F_{\pi_{\mathcal{G}-\mathcal{H}}}(A_{\mathcal{G}/\mathcal{H}})$. (Note that here I disregarded the possibility that $\mathcal{H}$ itself contains a subdiagram that is divergent; this will be handled later.) For the path vertex, the leading divergence of the subintegral is not identical to any in the percolation theory, as the sum of F_π over $\pi_\mathcal{H}$ does not reduce to a factor unity for the subdiagram; the 2-point vertex function with a mass (ϕ^2) insertion, which it resembles, is different,

though it has the same degree of divergence [an example of such an integral was discussed in eq. (3.90)]. Nonetheless, the weighted sum of F_π has similar factorization properties.

Motivated by these observations, I aim to prove that my perturbation expansion for the path vertex function can be renormalized in a manner very similar to that for percolation. Indeed, the mass, field, and coupling renormalizations will be exactly the same as in percolation, even when they occur inside a 1PI diagram for the path vertex function (we saw earlier that this is so for the self energy parts outside the 1PI path vertex function, that is connected to this function by a single line). For the path vertex itself, the renormalization works and takes a similar form as that for a mass insertion in percolation, but the coefficients are different. I will prove this to all orders in perturbation theory. First, I will establish that it is possible to perform subtractions as in ordinary field theory Feynman diagrams (e.g. for percolation), with the result that our subtracted amplitudes are non-diverging for each graph in every order in perturbation theory. Then I will show that, because of a cancellation of some sub-leading pieces involving self-energy insertions, the subtractions take the same form as in percolation, as indicated above. This then leads almost immediately to the RG equations, and the epsilon expansion for the exponents.

Renormalization of Feynman integrals

The idea for rendering the Feynman integral associated to a diagram finite is intuitively simple. We identify all the diverging sub-integrations associated to subdiagrams of the types already listed above (called "renormalization parts"), for which the divergence is related to the behavior of the integrand as a corresponding set of α parameters is scaled to zero, and then subtract away these parts of the integrand. One would hope that the resulting integrand is then convergent, and even absolutely convergent. It is necessary to prove this non-obvious result, which I will do using results from the literature.

The procedure is somewhat complex because a given diagram may contain several diverging subdiagrams. The subdiagrams may themselves contain diverging sub-subdiagram (these are revealed by considering several dilatation parameters ρ_i attached to distinct subdiagrams, which go to zero in some order). These possibilities cause no problems for disjoint subdiagrams (that

have no common edges), nor for a nested sub-subdiagram (entirely contained in a subdiagram). The case of subdiagrams that are neither disjoint nor nested, called "overlapping divergences", is more difficult, but turns out not to be a problem. One makes subtractions corresponding only to non-overlapping and nested subdiagrams. The procedure was defined by Bogoliubov and Parasiuk in recursive terms [109], finiteness was proved by Hepp [110], and a non-recursive definition in terms of "forests" was given by Zimmerman [111]. Together, this formulation is called the BPHZ method. These authors worked in terms of momentum-space integrals. For the later formulation and proofs within the parametric formulation, see Refs. [106, 108] and the review in Ref. [101].

Renormalization of parametric Feynman integrals

I briefly outline the result due to Bergère and Lam [106] that I will use. First it will be useful to introduce the "generalized Taylor operators" $\mathcal{T}^n$ [106, 108]. For a function $f(x)$ of a positive variable x that behaves as $f(x) \sim a_0 x^\nu$ as $x \to 0$ ($a_0 \neq 0$), such that $x^{-\nu} f(x)$ is infinitely-differentiable on $[0, a)$ ($a > 0$), and for our purposes with ν an integer, (such a function is said to have the Taylor series property) these are defined for any integer n to extract the Laurent-like series of terms:

$$\mathcal{T}^n f(x) = a_0 x^\nu + a_1 x^{\nu+1} + \ldots + a_{\nu+n} x^n, \tag{3.121}$$

(where a_0, $\ldots a_{\nu+n}$ are complex numbers) with properties $\mathcal{T}^n f(x) = 0$ if $n < \nu$, and $(1 - \mathcal{T}^n) f(x) \sim x^q$ with $q > n$. While the series has the Laurent form, I do not assume f is complex differentiable away from 0, and the coefficients can be calculated from f at positive x only, by ordinary Taylor expansion of $x^{-\nu} f(x)$ at $x = 0$. For a function of several variables $x_1, x_2, \ldots$, we may define generalized Taylor operators $\mathcal{T}^{n_i}_{x_i}$ similarly by acting with one of them at a time, but we must be careful as they do not generally commute.

In the following these operations will be applied acting on some subset E' of α's for a graph $\mathcal{G}$ by a dilatation parameter ρ as $\rho \to 0$, and then setting $\rho = 1$ in the result: $\mathcal{T}^n_{E'} f(A_\mathcal{G}) = \mathcal{T}^n_\rho f(\{\rho^2 \alpha_i : i \in E'\}, \{\alpha_i : i \notin E'\})|_{\rho=1}$. Thus these extract precisely the leading and subleading terms that we have been discussing, up to order n. Here when $f(\{\rho^2 \alpha_i : i \in E'\}, \{\alpha_i : i \notin E'\})$ has the Taylor series property as a function of ρ, we say it has it with respect to the set E', and it is

in this case that the operator $\mathcal{T}^n_{E'}$ is defined.

We will need some definitions for properties of the functions to which the Theorem applies. We will consider what Bergère and Lam [106] call a "nest" of edge subsets, which is a filtration, that is a set $\mathcal{N} = \{E_1, \ldots, E_r\}$ of edge subsets such that

$$\emptyset \subset E_1 \subset E_2 \subset \cdots \subset E_t \subseteq E(\mathcal{G}) \tag{3.122}$$

in which the inclusions are strict except possibly the last. For a function $Z(A_\mathcal{G})$, we say that it has the "simultaneous Taylor series property" with respect to the filtration $\mathcal{N}$ if there is a set of integers ν_{E_r} such that $(\prod_{r=1}^t \rho_r^{-\nu_{E_r}}) Z(\rho^2 A_\mathcal{G})$ has simultaneous Taylor series in the set of ρ_r near, and does not vanish at, $\rho_r = 0$ for all r; here $\rho^2 A_\mathcal{G}$ stands for the ordered set $A_\mathcal{G}$ of α's, but each α_i acquires a factor ρ_r^2 for each subset E_r to which i belongs. For example, the function $1/(\alpha_1 + \alpha_2)$ has the simultaneous Taylor series property for the filtration $E_1 = \{1\}$, $E_2 = \{1, 2\}$.

Now I can state (a special case of) the theorem of Bergère and Lam: if (i) $Z(A_\mathcal{G})$ is infinitely differentiable for $0 < \alpha_i < \infty$; (ii) $Z(A_\mathcal{G})$ and its α derivatives are polynomially bounded when arbitrary subsets of $A_\mathcal{G}$ are scaled to ∞; (iii) $Z(A_\mathcal{G})$ has the simultaneous Taylor series property with respect to every filtration $\mathcal{N}$ of edge subsets, then the integral

$$I_R = \int_0^\infty \prod_{i=1}^{|E(\mathcal{G})|} d\alpha_i \, e^{-\sum_i \alpha_i t} R[Z(A_\mathcal{G})] \tag{3.123}$$

with $t > 0$ is absolutely convergent. Here the R operation is the subtraction operator which can be defined as

$$R = 1 + \sum_\mathcal{N} \prod_{E' \in \mathcal{N}} (-\mathcal{T}_{E'}^{-2|E'|}), \tag{3.124}$$

where the sum is over all filtrations $\mathcal{N}$ of the set $A_\mathcal{G}$ of α's.

In its general form, the theorem applies to many integrals that are not related to Feynman diagrams in any obvious way. Now we wish to apply it to the Feynman integrals in my perturbation expansion. First I point out that these integrals do satisfy the hypotheses of the theorem. Indeed, the integrands of our integrals contain factors that occur in the field theory of percolation, which

for this purpose is no different from a cubic-interaction scalar field theory, times the factor F_π for some ordering π (times d_{MSF} and summed over π, but we need not consider this here; this sum can be exchanged with the integral and then taken under the R operation if desired). The integrand in the cubic theory satisfies the conditions, and it is easy to see that the F_π factor does not change this.

To go further, I note that when applied to Feynman integrals based on a graph (the graph made no appearance in the statement of the theorem), R can also be expressed in many other ways, one of which is as the sum over forests of renormalization parts [106]. As we know, a forest is a collection of trees, but here the trees are not spanning trees on our lattice or our graph $\mathcal{G}$. Instead, a forest is any set of renormalization parts in $\mathcal{G}$ (which are 1PI connected 2- or 3-point subdiagrams), such that for any two such parts in the set, either one is entirely contained in the other (both for its vertices and its edges), or else they are disjoint. (Often in the literature, a forest is pictured as a set of non-intersecting boxes overlaid on the depiction of the Feynman diagram.) In this form for R, the sum over all filtrations is replaced by a sum over all forests, and each edge set E' in the product is that of a single renormalization part belonging to that forest. I note that in the BPHZ formulation, whether in parametric form or not, no divergent integral or cutoff is mentioned. The subtractions are carried out instead on the integrand (which however, before the subtractions are performed, does have the property of diverging more strongly in some limits).

Renormalization of MSF path vertex integrals

I will apply the theorem to the integral for a diagram $\mathcal{G}$, in which the sum over orderings, and d_{MSF} factors, are taken into the integrand. That is,

$$Z(A_{\mathcal{G}}) = \sum_{\pi} d_{\text{MSF}}(\mathcal{G}|\pi) F_\pi(A_{\mathcal{G}}) P^{-d/2}(A_{\mathcal{G}}) \exp(-\mathbf{k}_{\text{ext}}^T \Delta^{-1} \mathbf{k}_{\text{ext}}). \tag{3.125}$$

My earlier remarks imply that for each renormalization part, the subtractions (in forest form) remove precisely all the superficially divergent pieces and no more. Thus in this form, the R operation is exactly what one might expect it to be from the discussion preceding the statement of the theorem, and the theorem says that these subtractions result in an absolutely convergent integral. For subdiagrams $\mathcal{H}$ of a diagram $\mathcal{G}$, these subtractions are exactly the same as those for a

diagram of the same type (number of external points, and presence or absence of the path vertex), with one exception. This is the subleading superficial divergence in the case of $\mathcal{H}$ a self-energy subdiagram. In the subleading generalized Taylor expansion (that is, $\mathcal{T}^{-2|E(\mathcal{H})|} - \mathcal{T}^{-2-2|E(\mathcal{H})|}$ acting on the dilatation factor for the subdiagram), part of it comes from expanding F_π to order α (for some α in the subdiagram) times the leading term from the percolation integrand; the terms from F_π arise from orderings in class $\pi_{[1]}$ and from subleading terms in class $\pi_{[0]}$. This does not correspond to the subtraction made to any whole diagram, and would thus be difficult to include in the renormalization scheme. Fortunately, these subtractions cancel, not for the given diagram, but between diagrams of the same order that differ only in the placement of the self-energy insertion in the graph. This cancellation result will be proved in section 3.4.6 below.

Hence because we are always interested in the sum of all diagrams in each order anyway, the only subtractions that have to be made correspond to those that would be made to $\mathcal{G}$ when it is a renormalization part. It follows that the subtractions correspond to subtracting the Taylor series in $\mathbf{k}^2$ for the subdiagram, where $\mathbf{k}$ is the wavevector entering the subdiagram, and replacing the original graph by the quotient by the subdiagram, times these Taylor coefficients in place of the subdiagram. For the vertex and path vertex cases, the subtraction is simply at zero wavevector, while for the self-energy the first order term in $\mathbf{k}^2$ has to be subtracted also. This is easily seen, as the subtractions to $\mathcal{G}$ itself are just its Taylor expansion in $\mathbf{k}^2$ to the given order [101, 106, 108]. If we include the zero-loop parts of the vertex functions, this implies that the renormalized vertex functions in this renormalization scheme obey

$$\begin{aligned}
\Gamma_R^{(2)}(\mathbf{0}; g, t) &= t, \\
\frac{d}{d\mathbf{k}^2}\Gamma_R^{(2)}(\mathbf{k}; g, t)\Big|_{\mathbf{k}=\mathbf{0}} &= 1, \\
\Gamma_R^{(3)}(\mathbf{0}, \mathbf{0}, \mathbf{0}; g, t) &= g, \\
\Gamma_R^{(2,1)}(\mathbf{0}, \mathbf{0}, \mathbf{0}; g, t) &= 1, \\
\Gamma_R^{(2,\mathrm{PV})}(\mathbf{0}, \mathbf{0}, \mathbf{0}; g, t) &= 1.
\end{aligned} \tag{3.126}$$

In view of the condition on the 2- and 3-point vertex functions at $\mathbf{k} = \mathbf{0}$, the coupling and mass-

squared appearing in the propagators in the expansion can be identified with the *renormalized* values, so there is no subscript zero on these quantities.

Now that the renormalized perturbation series defining the Γ_R are known to be finite, for example at non-zero wavevectors away from the point $\mathbf{0}$ at which the above conditions are given, we can modify the renormalization scheme. Namely, we can add a finite part (more accurately, a series of finite terms) to each subtracted piece in the definition of the renormalized integrand. These can be chosen in each order to modify the renormalization conditions, and the combinatorics again works out. This changes the renormalization scheme, and for example we can modify the conditions above to specify values at non-zero wavevectors (except for the mass-squared):

$$\begin{aligned} \Gamma_R^{(2)}(\mathbf{q}=0;g,t) &= t, \\ \frac{d}{d\mathbf{q}^2}\Gamma_R^{(2)}(\mathbf{q};g,t)\Big|_{|\mathbf{q}|=\kappa} &= 1, \\ \Gamma_R^{(3)}(\{\mathbf{q}_i\};g,t)\Big|_{\mathrm{SP}} &= g, \\ \Gamma_R^{(2,1)}(\{\mathbf{q}_i\};g,t)\Big|_{\mathrm{SP}} &= 1, \\ \Gamma_R^{(2,\mathrm{PV})}(\{\mathbf{q}_i\};g,t)\Big|_{\mathrm{SP}} &= 1. \end{aligned} \tag{3.127}$$

Here SP $=$ symmetry point denotes a symmetric configuration of external momenta $\mathbf{q}_1$, $\mathbf{q}_2$, $\mathbf{q}_3$, which (by rotational symmetry) we take to be any triple satisfying $\mathbf{q}_i^2 = \kappa^2$ for $i = 1, 2, 3$, $\mathbf{q}_i \cdot \mathbf{q}_j = -\kappa^2/2$ ($i \neq j$). Note that g and t now have a different meaning than before. In this form, we can now set $t = 0$ and work directly at the critical point, as in each order in perturbation theory the non-zero wavevector scale κ prevents the left-hand-sides from diverging in the infrared (the self-energy $-\Gamma^{(2)}(\mathbf{0};g,0)$ is not infrared divergent). This renormalization at zero mass-squared is quite convenient technically.

As I mentioned above, the BPHZ subtraction scheme requires no reference to, nor use of, a cutoff. It is possible to develop the RG equations directly from this scheme, working with non-divergent expressions only, and leading for example to the Callan-Symanzik equation when the renormalization scheme at zero wavevector, non-zero t is used [112]. However, for calculational purposes, I prefer to write intermediate quantities in terms of expressions that diverge as $\Lambda \to \infty$

as in traditional approaches. The bare vertex functions are given by the original, unsubtracted Feynman integrals with cut-off, including as always the $d_{\rm MSF}F_\pi$ factors. For emphasis, I now write these as Γ_0's. They are viewed as functions of the bare coupling g_0 and mass-squared t_0, as well as the wavevectors and cutoff λ. Then all the subtractions that define the renormalized amplitudes can be collected into changes of the parameters to g and $t = 0$, and changes in the scale of the "operators" ϕ, ϕ^2 and that described by the path vertex. That is

$$\begin{aligned} &Z_\phi^{N/2}(g_0,\kappa,\Lambda)Z_{\phi^2}^{L}(g_0,\kappa,\Lambda)\Gamma_0^{(N,L)}(\{\mathbf{q}_i\},\{\mathbf{q}_j\};g_0,t_0,\Lambda) \\ &\qquad = \Gamma_R^{(N,L)}(\{\mathbf{q}_i\},\{\mathbf{q}_j\};g,\kappa,\Lambda), \end{aligned} \tag{3.128}$$

$$\begin{aligned} &Z_\phi(g_0,\kappa,\Lambda)Z_{\rm PV}(g_0,\kappa,\Lambda)\Gamma_0^{(2,{\rm PV})}(\{\mathbf{q}_i\};g_0,t_0,\Lambda) \\ &\qquad = \Gamma_R^{(2,{\rm PV})}(\{\mathbf{q}_i\};g,\kappa,\Lambda), \end{aligned} \tag{3.129}$$

and in the limit $\Lambda \to \infty$ the dependence of all Γ_R on Λ drops out. These equations require five equations to define the dependence of Z_ϕ, Z_{ϕ^2}, $Z_{\rm PV}$, g_0, and t_0 on g, κ, and Λ, and these are provided by the five conditions (3.127), when these are expanded in perturbation theory in g_0. At this point the treatment of my theory has come to closely resemble an ordinary field theory, the main difference being the form of the Feynman rules for calculating $\Gamma^{(2,{\rm PV})}$. The most important conclusion of the analysis is that the path vertex is renormalized multiplicatively by $Z_{\rm PV}$.

3.4.6 Cancellation proof for subleading terms

In this section I present the proof that the particular subleading terms in the Laurent expansion as the α's in a self-energy part (not containing $\mathbf{z}$) go to zero, that do not appear for the self-energy in an external line, actually all cancel among graphs with the same self-energy part inserted in different edges.

First I show that the dependence on the α's in a self-energy of the weighted sum of F_π factors simplifies. We suppose throughout this section that we consider a fixed graph $\mathcal{G}_0$ with an ordering π_0, and we then modify this graph to obtain $\mathcal{G}_i$ by inserting a given self-energy graph $\mathcal{H}$ on an edge i of $\mathcal{G}_0$. Thus, the edge i is replaced by two edges i', i'', with the self-energy $\mathcal{H}$ in between. In the

parametric integral for the diagram, the parameter α_i in $\mathcal{G}_0$ is replaced by parameters $\alpha_{i'}$, $\alpha_{i''}$, and there are additional parameters for the edges of $\mathcal{H}$.

We know from section 3.4.4 that the d_{MSF} weight for $\mathcal{G}_i$ is determined by the ordering $\pi_{\mathcal{G}_i/\mathcal{H}}$ in which the cost replacing the original ℓ_i is the largest of $\ell_{i'}$, $\ell_{i''}$ and those in $\mathcal{H}$, independent of how these are ordered relative to each other. Moreover, the weight factors as

$$d_{\text{MSF}}(\mathcal{G}_i|\pi) = d_{\text{MSF}}(\mathcal{H}) d_{\text{MSF}}(\mathcal{G}_i/\mathcal{H}|\pi_{\mathcal{G}_i/\mathcal{H}}). \tag{3.130}$$

Throughout the argument, we will compare cases in which $\pi_{\mathcal{G}_i/\mathcal{H}} = \pi_0$ (in an obvious sense), and is fixed. The sum of the F_π factors over the orderings of the edges that replace i can be calculated, and this is done most easily by returning to the original calculation of F_π from the action of the $\mathcal{O}_{\text{MSF}}$ operator in section 3.4.2. The desired sum has the effect of simplifying the integro-differential operator to the following form, and acting under the parametric integral gives

$$\int_D \cdots \int \prod_{\epsilon \in E(\mathcal{G}_0)} dt_\epsilon \, \frac{d}{dt_\epsilon} e^{-\sum_{\epsilon \in E(\mathcal{G}_i)} \alpha_\epsilon t_\epsilon} = \prod_{j=1}^{E(\mathcal{G}_0)} \frac{\alpha'_{\pi_0(j)}}{\sum_{k=1}^{j} \alpha'_{\pi_0(k)}}, \tag{3.131}$$

where (i) D is the usual ($|E(\mathcal{G}_0)|$-dimensional) integration domain for $\mathcal{G}_0$ with ordering π_0, defined by $t_{\pi_0(1)} > \cdots > t_{\pi_0(|E(\mathcal{G}_0)|} > t_0$, and in the integrand, $t_{i'}$, $t_{i''}$, and the t's associated to edges in $\mathcal{H}$ are all set equal to t_i, and (ii) α'_j are the same as α_j except for α'_i, which is the sum of $\alpha_{i'}$, $\alpha_{i''}$ and all the α's in $\mathcal{H}$. The product on the right-hand side is simply F_{π_0} for $\mathcal{G}_0$ but with this substitution; I denote it $F'_{\pi_0}(i)$.

A second trick that is commonly used for parametric integrals is also useful: if the integrand only depends on the sum of two parameters, say α_1 and α_2, then these integrations can be combined into a single integral over α which takes the place of $\alpha_1 + \alpha_2$, at the cost of introducing a factor α into the integrand:

$$\int d\alpha_1 d\alpha_2 \, \ldots = \int d\alpha \, \alpha \, \ldots . \tag{3.132}$$

(This can be generalized to any number of variables, but I do not require that.) It can be shown that, like the weighted sum of F_π factors, the rest of the parametric integrand only depends on the

sum $\alpha_{i'} + \alpha_{i''}$ (this can be shown by some further use of the relation, mentioned earlier, of this integrand to combinatorics of weighted spanning trees, which I do not enter into). Then we use this result to recover an integral over a single α_i in place of those two.

I now consider the generalized Taylor expansion of the integrand with respect to a dilatation parameter ρ applied to the α's in $\mathcal{H}$. The leading behavior is seen to give simply the integrand for $\mathcal{H}$ from percolation (with zero external wavevectors), times the integrand and F_{π_0} factor for the quotient graph, as discussed above. I now turn to the subleading terms of a particular form: those that come from expanding the above factor $F'_{\pi_0}(i)$ to first order in ρ^2 (or simply in α's in $\mathcal{H}$), times the leading behavior of the rest of the integrand, as $\rho \to 0$. The rest of the integrand factors into that for the quotient graph times that for the subgraph, and we make use of the technique for replacing the $\alpha_{i'}$, $\alpha_{i''}$ by α_i. This factor in the integrand is now independent of which edge i of $\mathcal{G}_0$ was chosen for insertion of $\mathcal{H}$. The d_{MSF} factor is also independent of i because of the choices of ordering made earlier. Let us write $\sum \alpha_{\text{s.e.}}$ for the sum of α's associated to $\mathcal{H}$. There are two types of terms in the expansion of $F'_{\pi_0}(i)$ at first order in $\sum \alpha_{\text{s.e.}}$ for each i: a) those in which $\sum \alpha_{\text{s.e.}}$ appears in the numerator, in which case it replaces α_i, so giving the factor $\sum \alpha_{\text{s.e.}}/\alpha_i$ times F_{π_0}; b) those in which it comes from expanding a denominator, which must be one of those indexed $j \geq \pi_0^{-1}(i)$. This gives a factor $-\sum \alpha_{\text{s.e.}}/\sum_{k=1}^{j} \alpha_{\pi_0(k)}$ times F_{π_0}. We remember to multiply by α_i (because we replaced two α's by this one), and then sum over the positions i of the self-energy insertion. This gives

$$\begin{aligned} \sum_i 1 - \sum_i \sum_{j: j \geq \pi_0^{-1}(i)} \frac{\alpha_i}{\sum_{k=1}^{j} \alpha_{\pi_0(k)}} &= \sum_i 1 - \sum_j \frac{\sum_{i: \pi_0^{-1}(i) \leq j} \alpha_i}{\sum_{k=1}^{j} \alpha_{\pi_0(k)}} \\ &= 0, \end{aligned} \tag{3.133}$$

times other common factors. This completes the proof.

What I have shown is that this type of subleading term actually cancels in the sum of diagrams in each order. Because the superficial degree of divergence of a self-energy diagram at six dimensions is 2, these subleading parts are also superficially divergent, and are subtracted for each diagram by the R operation defined earlier. We view the present result as showing that the

subtracted terms cancel, and because the sum of the subtracted integrals is finite, these subtracted terms can be dropped and the result is still finite. The cancellation is independent of other subtractions associated with renormalization parts, which might appear either inside $\mathcal{H}$ or disjoint from it. In particular, this allows the cancellation to be made for any number of self-energy insertions in a graph. The remaining terms subtracted from a self-energy insertion are then exactly those that occur in percolation, for any number of self-energy insertions.

The principle underlying these pleasant cancelations is not entirely clear to us. It seems likely (because they involve the derivative of the self-energy in percolation with respect to t) that they are associated with the notion that the $\mathcal{O}_{\rm MSF}$ operator should be renormalized, so that it acts on t rather than t_0. This operator has the property that it is invariant under any reparameterization of the variables $t_\epsilon \to T_\epsilon = T(t_\epsilon)$ provided that T is a monotonic function and has no explicit dependence on ϵ (this invariance of the geometry of MSTs is related to that emphasized in Ref. [10]). This property of $\mathcal{O}_{\rm MSF}$ was not explicitly used anywhere in my construction. I will not attempt to give here a conceptual proof using these ideas.

3.5 One-loop renormalization group calculation

I have proved in the preceding section that the diagrammatic expansion for the MSF path theory is renormalizable. These proofs were technical, but having established this fact, we are free to make use of standard RG methods such as those discussed in Ref. [102]. I continue to take a rather pedagogical approach in this section.

As I explained in sections 3.3.4 and 3.4, renormalizability of a theory means that we may absorb the strong Λ-dependence of all correlation functions into a finite number of parameters and the overall scale of the correlations, at the cost of introducing another scale κ. In the scheme I use, in which the renormalization conditions are at zero renormalized mass-squared $t = 0$ and non-zero wavevector of order κ, the precise statement is that functions g, Z_ϕ, Z_{ϕ^2}, $Z_{\rm PV}$ exist such that (here I append subscripts 0 to denote the "bare" vertex functions as constructed above, with cutoff λ)

$$Z_\phi^{N/2}(g_0,\kappa,\Lambda)Z_{\phi^2}^{L}(g_0,\kappa,\Lambda)\Gamma_0^{(N,L)}(\{\mathbf{q}_i\},\{\mathbf{q}_j\};g_0,t_0,\Lambda) = \Gamma_R^{(N,L)}(\{\mathbf{q}_i\},\{\mathbf{q}_j\};g,\kappa), \quad (3.134)$$

$$Z_\phi(g_0,\kappa,\Lambda)Z_{\rm PV}(g_0,\kappa,\Lambda)\Gamma_0^{(2,{\rm PV})}(\{\mathbf{q}_i\};g_0,t_0,\Lambda)=\Gamma_R^{(2,{\rm PV})}(\{\mathbf{q}_i\};g,\kappa), \tag{3.135}$$

where Γ_R are independent of Λ as $\Lambda\to\infty$ with g, $t=0$ fixed, up to corrections vanishing in this limit. (Here L is the number of insertions of ϕ^2, and $\{\mathbf{q}_j\}$ is the set of corresponding wavevectors.) I will also now introduce dimensionless versions of the bare and renormalized couplings g_0, g:

$$u_0^2\equiv\frac{g_0^2\kappa^{-\varepsilon}}{(4\pi)^{d/2}},\qquad u^2\equiv\frac{g^2\kappa^{-\varepsilon}}{(4\pi)^{d/2}}. \tag{3.136}$$

I introduced an angular factor $(4\pi)^{d/2}$ in the above definitions for later convenience to simplify expressions. For the vertex functions not containing the path vertex, the functions and their renormalization is exactly as percolation, and this is also true of the following calculations; I include some details anyway to provide checks on the calculation.

The RG equations are obtained from the observation that the bare functions Γ_0 are independent of κ when written in terms of g_0, so $\kappa\partial\Gamma_0/\partial\kappa=0$ at fixed g_0, t_0, Λ. Using the definition of Γ_R we obtain

$$\left(\kappa\frac{\partial}{\partial\kappa}+\beta(u)\frac{\partial}{\partial u}-\frac{N}{2}\gamma_\phi(u)+L\gamma_{\phi^2}(u)\right)\Gamma_R^{(N,L)}(\{\mathbf{q}_i\},\{\mathbf{q}_j\},u,\kappa)=0, \tag{3.137}$$

and

$$\left(\kappa\frac{\partial}{\partial\kappa}+\beta(u)\frac{\partial}{\partial u}-\gamma_\phi(u)+\gamma_{\rm PV}(u)\right)\Gamma_R^{(2,{\rm PV})}(\{\mathbf{q}_i\},u,\kappa)=0 \tag{3.138}$$

In each of these equations the first and second partial derivatives are at fixed u and fixed κ, respectively. The RG β and γ functions appearing in equations (3.137), (3.138) are defined as

$$\beta(u) = \kappa\frac{\partial u}{\partial\kappa}\bigg|_{g_0,\Lambda}, \tag{3.139}$$

$$\gamma_\phi(u) = \kappa\frac{\partial\log Z_\phi}{\partial\kappa}\bigg|_{g_0,\Lambda}, \tag{3.140}$$

$$\gamma_{\phi^2}(u) = -\kappa\frac{\partial\log Z_{\phi^2}}{\partial\kappa}\bigg|_{g_0,\Lambda}, \tag{3.141}$$

$$\gamma_{\rm PV}(u) = -\kappa\frac{\partial\log Z_{\rm PV}}{\partial\kappa}\bigg|_{g_0,\Lambda}, \tag{3.142}$$

and are finite as $\Lambda \to \infty$ [102]. Hence in the limit they are independent of Λ, and so also of κ, because u, β, and all γ's are dimensionless; they are simply power series in u.

I then impose the following renormalization conditions, which are those we reached in section 3.4, but written now with $t = 0$ and u in place of g. These serve to fix the dependence of the parameters on one another:

$$\begin{aligned}
\Gamma_R^{(2)}(\mathbf{q}=0,u,\kappa) &= t = 0,\\
\frac{d}{d\mathbf{q}^2}\Gamma_R^{(2)}(\mathbf{q},u,\kappa)\Big|_{|\mathbf{q}|=\kappa} &= 1,\\
\Gamma_R^{(3)}(\{\mathbf{q}_i\},u,\kappa)\Big|_{\mathrm{SP}} &= g,\\
\Gamma_R^{(2,1)}(\{\mathbf{q}_i\},u,\kappa)\Big|_{\mathrm{SP}} &= 1,\\
\Gamma_R^{(2,\mathrm{PV})}(\{\mathbf{q}_i\},u,\kappa)\Big|_{\mathrm{SP}} &= 1.
\end{aligned} \tag{3.143}$$

Here SP denotes a symmetry point of the external momenta $\mathbf{q}_1$, $\mathbf{q}_2$, $\mathbf{q}_3$, at which $|\mathbf{q}_i|^2 = \kappa^2$, as defined in section 3.4 (though the precise definition is unimportant). These conditions are now used to determine β and the γ's from the perturbation theory expansion in g_0 of the 1PI vertex functions Γ_0 with a fixed cutoff Λ. The expressions make sense provided g_0 is sufficiently small. I will calculate to one-loop order, which means that only the one-loop diagrams for the renormalization parts need to be calculated. This will give results for exponents to first order in $\varepsilon = 6 - d > 0$. (More generally, computing to $\mathcal{O}(\varepsilon^{\mathcal{L}})$ in the ε-expansion requires computing all the renormalization parts with $\mathcal{L}$ or fewer loops.)

Then the instances of equation (3.135) with which we need to deal are, to $\mathcal{O}(g_0^2)$,

$$\begin{aligned}
\Gamma_R^{(2)}(\mathbf{q},u,\kappa) &= Z_\phi\left((\mathbf{q}^2+t_0)-\Sigma(\mathbf{q})\right),\\
\Sigma(\mathbf{q},g_0) &= d_2 g_0^2 I_2(\mathbf{q});\\
\Gamma_R^{(3)}(u,\kappa)|_{\mathrm{SP}} &= Z_\phi^{3/2}\left(g_0 + d_3 g_0^3 I_3\big|_{\mathrm{SP}}\right);\\
\Gamma_R^{(2,1)}(u,\kappa)|_{\mathrm{SP}} &= Z_\phi Z_{\phi^2}\left(1 + d_2 g_0^2 I_3\big|_{\mathrm{SP}}\right);\\
\Gamma_R^{(2,\mathrm{PV})}(u,\kappa)|_{\mathrm{SP}} &= Z_\phi Z_{\mathrm{PV}}\left(1 + g_0^2 \sum_{\pi\in S_3} d_{\mathrm{MSF}}(\triangle|\pi)\mathcal{O}_{\mathrm{MSF}}(\pi) I_3(\pi)|_{\mathrm{SP}}\right).
\end{aligned} \tag{3.144}$$

Here d_2, d_3 are the percolation d-weights for these one-loop diagrams, $d_2 = -1$, $d_3 = -2$, and $d_{\rm PV}$ will be evaluated in a moment from the weights $d(\mathcal{G})$ and the operator $d_{\rm MSF}\mathcal{O}_{\rm MSF}$; the values will be substituted only at the end of the calculation. The negative sign in the equation for $\Gamma_R^{(2)}$ arises because of Dyson's equation.

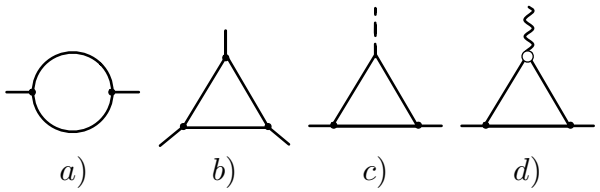

Figure 3.5: The 1PI one-loop diagrams for vertex functions that are ultraviolet divergent in six dimensions: a) self-energy, b) cubic interaction vertex function, c) mass-insertion vertex function, d) MSF path vertex function.

In equation (3.144), I_2 and I_3 are Feynman loop integrals which I now evaluate. To leading non-trivial order, we can evaluate them at six dimensions, retaining only the terms that diverge quadratically or logarithmically as $\Lambda \to \infty$. These terms may be extracted by any of the standard techniques for evaluating Feynman integrals, including those in section 3.4; I refer the reader to [102, 113, 114] in particular. Denoting this approximation by $\simeq$, we find

$$\begin{aligned} I_2(\mathbf{q}) &= \frac{1}{2}\int^{\Lambda} \frac{d^d\mathbf{k}}{(2\pi)^d}\,\frac{1}{\mathbf{k}^2(\mathbf{k}+\mathbf{q})^2}, \\ &\simeq \frac{1}{2}\frac{1}{(4\pi)^{d/2}}\left(\Lambda^2 - \frac{\mathbf{q}^2}{6}\log\frac{\Lambda^2}{\mathbf{q}^2}\right). \end{aligned} \tag{3.145}$$

The factor of $1/2$ appearing in $I_2(\mathbf{q})$ is a diagrammatic symmetry factor, $\mathcal{A} = 2$ (in the other diagrams, $\mathcal{A}$ or $\mathcal{A}' = 1$). Note that the bare propagators appearing in these integrals should have mass-squared $t_0 = t_c$ determined so that $t = 0$, but because t_c is $\mathcal{O}(g_0^2)$ we may consistently neglect its presence here. The other integral is

$$\begin{aligned} I_3|_{\rm SP} &= \int^{\Lambda} \frac{d^d\mathbf{k}}{(2\pi)^d}\,\frac{1}{\mathbf{k}^2(\mathbf{k}+\mathbf{q}_1)^2(\mathbf{k}-\mathbf{q}_2)^2}\bigg|_{SP}, \\ &\simeq \frac{1}{2}\frac{1}{(4\pi)^{d/2}}\log\frac{\Lambda^2}{\kappa^2}. \end{aligned} \tag{3.146}$$

For the path vertex function, as we saw above, $\mathcal{O}_{\rm MSF}(\pi, t_c)I_3|_{\rm SP}$ is independent of the ordering

π. (Unfortunately, this property does not hold to higher orders in perturbation theory for $\Gamma^{(2,\mathrm{PV})}$.) The result for a single ordering is

$$\mathcal{O}_{\mathrm{MSF}}(\pi, t_c) I_3|_{\mathrm{SP}} = \frac{1}{3!} I_3|_{\mathrm{SP}}. \tag{3.147}$$

Then for brevity I define d_{PV} such that

$$\sum_{\pi \in S_3} d_{\mathrm{MSF}}(\triangle|\pi) \mathcal{O}_{\mathrm{MSF}}(\pi, t_c) I_3|_{\mathrm{SP}} = d_{\mathrm{PV}} I_3|_{\mathrm{SP}}. \tag{3.148}$$

Hence $d_{\mathrm{PV}} = \frac{1}{6} \sum_\pi d_{\mathrm{MSF}}(\triangle|\pi) = -2/3$, as we saw above.

We may now solve (3.144) by requiring that the renormalized proper vertices defined on the right-hand side of (3.144) satisfy the normalization conditions (3.143) up to terms of $\mathcal{O}(u_0^4)$. We obtain to this order

$$\begin{aligned}
t_c &= \tfrac{1}{2} d_2 u_0^2 \Lambda^2, && (3.149)\\
u &= \kappa^{-\varepsilon/2} g_0 \left(1 + (d_3 - \tfrac{1}{4} d_2) u_0^2 \log \frac{\Lambda}{\kappa}\right), && (3.150)\\
Z_\phi &= 1 - \tfrac{1}{6} d_2 u_0^2 \log \frac{\Lambda}{\kappa}, && (3.151)\\
Z_{\mathrm{PV}} &= 1 + (\tfrac{1}{6} d_2 - d_{\mathrm{PV}}) u_0^2 \log \frac{\Lambda}{\kappa}. && (3.152)
\end{aligned}$$

Finally, to this order Z_{ϕ^2} is the same as Z_{PV} except that d_2 replaces d_{PV}. Note that the Z's are functions only of the dimensionless variables u_0 and κ/Λ, and we can set $d = 6$ (so $\varepsilon = 0$ and $g_0^2 = u_0^2/(4\pi)^3$)) in the one-loop terms, but not in zero-loop terms.

The RG functions are

$$\begin{aligned}
\beta(u) &= -\tfrac{1}{2}\varepsilon u + \left(\tfrac{1}{4} d_2 - d_3\right) u^3 + \mathcal{O}(u^5, \varepsilon u^3), && (3.153)\\
\gamma_\phi(u) &= \tfrac{1}{6} d_2 u^2 + \mathcal{O}(u^4, \varepsilon u^2), && (3.154)\\
\gamma_{\mathrm{PV}}(u) &= \left(\tfrac{1}{6} d_2 - d_{\mathrm{PV}}\right) u^2 + \mathcal{O}(u^4, \varepsilon u^2). && (3.155)
\end{aligned}$$

The fixed points of the RG are the values of $u = u^*$ at which $\beta(u) = 0$. Clearly, one fixed point

is at $u = 0$, but is unstable to the introduction of the cubic coupling for $d < 6$, because u grows as κ *decreases*, corresponding to the behavior at *larger* length scales. In six dimensions, u approaches zero logarithmically as κ decreases, because the coefficient of u^3 is $d_2/4 - d_3 = 7/4$ which is positive. Below six dimensions, there is another zero of β which results from the competition between the two terms in β, at

$$(u^*)^2 = \frac{2\varepsilon}{d_2 - 4d_3} + \mathcal{O}(\varepsilon^2). \tag{3.156}$$

Note that $(u^*)^2$ is positive for $d < 6$ ($\varepsilon > 0$). For $d > 6$, this fixed point is not relevant to percolation or MSTs.

The values of the γ's at the fixed point value of u give the "anomalous dimensions" of the various operators (except that in the case of ϕ, the anomalous dimension is $\gamma_\phi/2$). These are the difference of the total dimensions x of the operators from their canonical or engineering dimensions, which are the zero-loop values discussed earlier [102]. The co-dimension $d - x$ gives the fractal dimension of the associated geometric object (set of points). The most interesting dimension for us is that of the path vertex, which is $x_{\rm PV} = d - 2 + \gamma_{\rm PV}$. The codimension yields the fractal dimension $D_{\rm p} = 2 - \gamma_{\rm PV}$ of the path on the MSF(p_c), as discussed in section 3.3.3. We find

$$\gamma_{\rm PV}(u^*) \equiv 2 - D_{\rm p} = \frac{(d_2 - 6d_{\rm PV})\varepsilon}{3(d_2 - 4d_3)} + \mathcal{O}(\varepsilon^2), \tag{3.157}$$

that is

$$D_{\rm p} = 2 - \frac{\varepsilon}{7} + \mathcal{O}(\varepsilon^2). \tag{3.158}$$

This is the main quantitative result of this chapter.

The other γ's produce fractal dimensions related to properties of percolation, which also apply to MSF(p). First,

$$\begin{aligned} \gamma_\phi(u^*) \equiv \eta &= \frac{d_2\varepsilon}{3(d_2 - 4d_3)} + \mathcal{O}(\varepsilon^2) \\ &= -\frac{\varepsilon}{21} + \mathcal{O}(\varepsilon^2). \end{aligned} \tag{3.159}$$

Hence for the fractal dimension $D_{\rm perc}$ of the critical percolation clusters, which is $D_{\rm perc} = d - x_\phi$, we find

$$D_{\rm perc} = (d + 2 - \eta)/2 = 4 - \frac{10\varepsilon}{21} + \mathcal{O}(\varepsilon^2). \tag{3.160}$$

The other γ is γ_{ϕ^2}, for which the value at the fixed point can be obtained from the formula for $\gamma_{\rm PV}$ by replacing $d_{\rm PV}$ by d_2, that is

$$\gamma_{\phi^2}(u^*) \equiv 2 - D_{\rm sc} = \frac{5\varepsilon}{21} + \mathcal{O}(\varepsilon^2), \tag{3.161}$$

and so for the fractal dimension of the set of singly-connected edges (see section 3.3.3) we find

$$D_{\rm sc} = \nu^{-1} = 2 - \frac{5\varepsilon}{21} + \mathcal{O}(\varepsilon^2). \tag{3.162}$$

The values I have obtained for both exponents η and ν agree with those in the literature on percolation, to order ε [96, 97, 99], which provides a check on my calculation.

The comparison of $D_{\rm p}$ with $D_{\rm sc}$ raises some questions of inequalities obeyed by $D_{\rm p}$. There are also some other fractal dimensions defined for paths on critical percolation clusters which have been studied. These include $D_{\rm min}$, the fractal dimension of the shortest path on the cluster between the given points, and $D_{\rm max}$, the fractal dimension of the longest (self-avoiding) path between them [39]. Here the length of the path is the number of edges of the lattice that it traverses. Then the inequalities are fairly obvious: first, because all these paths must pass through the singly-connected edges, $D_{\rm sc}$ is the smallest of all, and the remaining inequalities

$$D_{\rm sc} \le D_{\rm min} \le D_{\rm p} \le D_{\rm max}, \tag{3.163}$$

follow from the definitions. To order ε, one has $D_{\rm min} = 2 - \varepsilon/6$ and $D_{\rm max} = 2 - \varepsilon/42$ [115], and all the inequalities are obeyed *strictly* by the results to this order. $D_{\rm p}$ is close but not equal to $D_{\rm min}$.

3.6 Postscript: remarks for $p > p_c$

The astute reader will notice that I have not really addressed the problem I set out to solve: my goal was a computation of the fractal dimension of the path on the MST between two well-separated points, and it is not obvious that I am justified in working at p_c: we really want information outside the critical scaling region, *i.e.* the dimensions of paths on scales $L \gg \xi$ with $p > p_c$. In this section I present a heuristic argument that the fractal dimension should, in fact, be equivalent to the quantity I calculated.

The expansion for MSF(p) that I have obtained is naturally organized as a low-density expansion, that is as an expansion in powers of p. For large graphs $\widehat{G}$, it becomes unwieldy, especially for p greater than around the percolation threshold (to the extent that a threshold can be associated with a finite graph, for example for a portion of a hypercubic lattice we can consider the threshold p_c of the infinite lattice). The corresponding expansion for the connectedness function $C_{\mathbf{x},\mathbf{y}}(p)$ in percolation must produce the answer 1 when $p \to 1$, but in a very complicated way, as a sum of a large number of terms. For the MST, obtained from MSF(p) as $p \to 1$, the probability that the path on the MST from $\mathbf{x}$ to $\mathbf{y}$ passes through $\mathbf{z}$ remains non-trivial in the limit, and is again given by a complicated set of terms. A general analysis of this sum on a large graph for $p > p_c$ would require a resummation of terms to allow for the presence of the "giant cluster" in the corresponding percolation. In the Potts model formulation of percolation, this is done by giving an expectation value to the Potts spin.

3.6.1 The "superhighways" conjecture

There are claims (the "superhighways" argument) [20] (see also [116]) that in the Kruskal process (on various families of graphs) the properties of the paths are essentially determined at the percolation threshold. In the terminology of [20], the "superhighway" is the incipient infinite cluster at p_c, which connects "roads" consisting of paths on finite clusters that get connected to the spanning cluster as p is raised above p_c. Numerically, the fraction of steps on a path on the MST that are already present at $p = p_c$, averaged over all pairs of end vertices for the path, goes to a constant [20]. This suggests that $D'_{\rm p} = D_{\rm p}$, where the prime distinguishes critical exponents as p_c

is approached from above rather than below. For the Bethe lattice and for Euclidean systems with $d > 6$ we can argue that both dimensions equal 2 (section 2.2.7), so equality holds, however, it is less clear whether it holds for $d < 6$.

Some support for $D'_{\rm p} = D_{\rm p}$ can be obtained from scaling arguments in percolation. Consider a hypercube of side L, and for bond percolation ask for the probability that two given opposite faces are connected when the occupation probability for each edge is p. The percolation threshold is the value of p above which the connection occurs with probability one as the size $L \to \infty$. For L finite, connection occurs at $p = p_c$ with low probability, but it occurs with probability approaching 1 at $p - p_c$ of order $1/L^{1/\nu_{\rm perc}}$, where $\nu_{\rm perc}$ is the correlation length exponent in percolation [39]. At this value of p, the correlation length ξ is of order $(p_c - p)^{-\nu_{\rm perc}} = \mathcal{O}(L)$, and scaling properties at scales less than ξ should be the same as those at threshold $p = p_c$. Hence we expect the fractal dimension $D'_{\rm p}$ of the path on the MST connecting the two faces (which is one of the paths on the percolation cluster that do so when the faces are first connected as p increases) to be $D_{\rm p}$. However, this path does not have the boundary conditions we wanted, as we only asked for the connection of the two faces, not of two given vertices in the interior of a system. When the separation R of the vertices is large, we may surround each by a nested sequence of concentric spheres of radii say $2^{-j}R/3$ for $j = 0, 1, \ldots$ and ask the same question for each annulus bounded by two of these spheres. Then the relevant p will be different for each sphere (higher p is required to make the connections to the vertices on smaller scales), but also the scaling holds for each one. In addition, the clusters that connect each pair of spheres must also become connected together to form a single cluster and an MSF path. Similar arguments apply to all of these. This does suggest, heuristically, that $D'_{\rm p} = D_{\rm p}$.

On examining this argument, a key part of it can be seen to be the idea that there is a unique candidate "superhighway" (critical percolation cluster) that is used to make connections over large distances on each scale, and so it is clear which ones must be connected by higher-cost "roads". This is a property of critical percolation clusters that holds for $d < 6$, but not for $d > 6$, where the number of large clusters visible in a window of size W is W^{d-6} [39, 46]. This behavior may itself underlie the result from chapter 2 that the number of connected components of the MST (or

of MSF(p) for $p > p_c$) has this same form W^{d-6} for $d \geq 6$ (but order one for $d < 6$).

3.7 Conclusions

The results of this chapter fall into three main parts. First, I constructed an exact expansion for the Kruskal process, or spanning forest MSF(p), in a series in powers of p, which terminates for a finite graph, and is analogous to a low-density expansion for percolation, or a high-temperature expansion in a statistical mechanical model. The expansion is for the probability that the path on the MSF from $\mathbf{x}$ to $\mathbf{y}$ passes through a vertex $\mathbf{z}$. Second, this expansion was used to obtain a continuum formulation (with cut-off) for $p \leq p_c$ (where p_c is the percolation threshold) in terms of Feynman diagrams (the region $p > p_c$ presents additional technical problems, and I will not discuss these further here). This expansion was then shown to be renormalizable to all orders in perturbation theory, so that the limit of infinite momentum-space cutoff (or zero lattice spacing) can be taken. Third, the renormalized perturbation expansion was used to calculate the fractal dimension of any path on MSF(p) at $p = p_c$, to first order in $\varepsilon = 6 - d$, for $d \leq 6$: $D_{\rm p} = 2 - \varepsilon/7 + \mathcal{O}(\varepsilon^2)$. For $d > 6$, $D_{\rm p} = 2$. If the "superhighways" idea is correct, then the same $D_{\rm p}$ also applies to the region $p > p_c$, in which we expect the path dimension to be independent of p on large enough length scales.

It is important to realize that it was by no means obvious at the outset that such a field-theoretic renormalization process would be possible. The problem is not obviously given by a local field theory, and my expansion is not based on an action principle (at least, not in any apparent way). Optimization is generally a non-local process as it involves making comparisons among (sums of) costs globally; however, this is also true when one wishes to minimize a Hamiltonian, even if its parameters (corresponding to costs) multiply local interaction terms. For minimum spanning trees, the definition of the allowed or "feasible" configurations (i.e. spanning trees) is not local either. It was not obvious that the expansion would be renormalizable with the same techniques employed in local field theory. Indeed, in the end my procedure worked thanks to unexpected and non-local cancellations of some subleading divergences (see section 3.4.6), for which I am unaware of any analogues in local field theories. Undoubtedly the underlying reasons for this success with

MSTs should be found in the applicability of Kruskal's greedy algorithm and its connection with percolation.

Future directions

The calculations can be extended in various ways. The exponents can be calculated to higher orders in ε, with increasing effort required for each additional order. The path vertex function, and not only its scaling dimension, can in principle also be studied, as can more general correlation functions with path vertices and mass-insertion vertices. In six dimensions, there are logarithmic corrections to the simple scaling with $D_{\rm p} = 2$ that holds for dimensions bigger than six, and these are calculable.

Independently of these applications of the renormalized perturbation expansion, the exact lattice low-density expansion could be studied in low orders (say, the first thirty terms) in any dimension d, as is conventionally done with high-temperature series. This would provide another way to obtain scaling dimensions for correlation functions. Such techniques are frequently very accurate.

A further question is the Borel summability of the perturbation expansion, or of the ε expansion for the exponents. If an asymptotic expansion of a function is Borel summable, then it uniquely determines that function [117]. If a few terms of the expansion are available, and it is believed to be Borel summable, then an improved estimate for the quantity of interest, such as an exponent, for a non-zero value of the parameter (say $\varepsilon = 1$) can be made, and for critical exponents these values may be very accurate (comparable with high-temperature series methods). For percolation at threshold, the asymptotic high-order behavior of the perturbation expansion has been shown in Ref. [118] to have the form that is a necessary condition for the expansion to be Borel summable. These results also apply to my theory, but again I also would need a similar result for the path vertex function. It would be interesting to find a technique to estimate the high order behavior of my expansion.

In conclusion, the introduction of the Kruskal process and geometric object MSF(p), based on an optimization problem, provides a rich area for study not unlike conventional critical phenomena. At $p = p_c$, many techniques can be applied to it. It illuminates numerical work on such problems

as optimal paths and transport in random media.

Chapter 4

Fractional quantum Hall wavefunctions

In this chapter I consider a completely different problem from those in the previous chapters: that of specifying possible phases of the fractional quantum Hall effect. Since the work of Laughlin [50], the main tool we have for describing such phases has been the formalism of trial wavefunctions, in which a microscopic, first-quantized electron wavefunction is given which is taken to be representative of the entire phase: we assume the actual, physical wavefunctions will be adiabatically connected to the trial wavefunctions, and the validity of the trial wavefunctions may be checked by comparison with numerical exact diagonalization studies.

In the literature there are two main ways to describe a trial wavefunction for the fractional quantum Hall effect: either as a block of some conformal field theory (CFT), or as the zero-energy eigenstate of a specially-chosen projection Hamiltonian. Both approaches have strengths and weaknesses: with the former definition, topologically-invariant defining characteristics of the phase such as braiding statistics of quasihole excitations (see section 4.1.4) may easily be read off from the field content of the CFT, while states defined using the latter definition are easier to realize explicitly and work with numerically. Ultimately, it would be desirable to have some tools for translating between these two methods, and this is the problem I attempt to address in this chapter.

The specific theories studied in this chapter were motivated by a recent result of Simon [52], who evaluated the conformal block corresponding to the amplitude of an arbitrary number of supercurrents in a superconformal field theory (SCFT) at arbitrary c (see section 4.3.4) and found that it

took a form similar to those encountered for clustered FQHE states obtained from several-particle projection Hamiltonians. Specifically, there are two possible ways the wavefunction can vanish as three particles are brought together (see section 4.3.1); the Hamiltonian projects out one direction in this two-dimensional space, and that direction is described by a free parameter which is related to the central charge c of the CFT.

Simon's result correctly describes the angular-momentum ground state for any c. In section 4.3.4 I construct all excited states allowed by the Hamiltonian and find that the resulting basis is, in fact, independent of c: although the basis functions themselves change, there are still the same number of linearly-independent functions at any level of angular momentum. This is confirmed by counting the states: the character thus obtained is the character of the CFT describing chiral edge excitations of the state, and for this basis the character is found to agree with that of a non-minimal SCFT at generic central charge.

Stable quantum Hall phases are described by unitary CFTs [53], and obtaining a unitary minimal theory requires excluding additional states in accordance with the singular vector structure of the SCFT. I am able to construct basis states and find the additional projection operator needed to obtain the (non-unitary) $SM(2,8)$ SCFT. The simplest unitary minimal SCFT appears to be the tricritical Ising model, $SM(3,5)$, but here I have been able to make less progress. The necessary projection operator here appears to require seven-particle interactions.

The new results of this chapter are the basis for translationally-invariant symmetric polynomials (4.71) (and additionally (4.76)) and proof of its linear independence, the basis of excited states for the $SM(2,8)$ minimal SCFT in (4.109), and the computational work on finite-N wavefunctions derived from the tricritical Ising model $SM(3,5)$ at the end of section 4.3.6, in particular the characters at finite particle number (4.136) and the explicit form of the level three singular vector (4.137).

Chapter outline

I begin by going through some basic elements of the theory behind the fractional quantum Hall effect in section 4.1, developing the trial wavefunction formalism in section 4.1.2 and describing the

most successful and well-studied trial wavefunctions, the Laughlin states, in section 4.1.3. These states are best thought of as going hand-in-hand with special projection Hamiltonians, constructed so that the states in question are the densest zero-energy eigenstates. In section 4.1.4 I elaborate more on the theoretical underpinnings connecting trial wavefunctions to topological field theory, motivating the choice of trial wavefunctions from among the blocks of CFTs which was pioneered in [119].

This necessitates introducing conformal field theory into the picture, and section 4.2 is devoted to an extremely compressed development of that subject. I discuss the consequences of conformal invariance in field theory at the classical level in section 4.2.1 and at the quantum level in section 4.2.2, arriving at the Virasoro minimal models in section 4.2.3. I come full circle and describe the relation between topological field theory in $2+1$ dimensions and conformal field theory in two dimensions underlying the description of the FQHE in section 4.2.4. Finally, in section 4.2.5 I briefly discuss the consequences of adding supersymmetry to conformal invariance, which is relevant for the specific project I undertake in the following section.

In section 4.3 I turn my attention to the investigation of the problem outlined above. I begin in section 4.3.1 by describing how projection Hamiltonians may be extended to several-body interactions, following [64]. In section 4.3.2 I discuss a simple argument (originally due to Read in the context of parafermionic FQHE states [65, 66]) which allows one to count and explicitly construct trial wavefunctions corresponding to all excitations of a clustered ground state obtained from a projection Hamiltonian. I refer to this as the "filtration argument," and use it to obtain the results of the following sections.

I begin the presentation of new results in section 4.3.3, where I present an alternate basis for the multi-body projection Hamiltonians which is better suited to use in filtration-type arguments. In section 4.3.4, I recapitulate relevant results in unpublished notes by Read [67], who uses a filtration argument to construct a basis for the zero-energy eigenspace projection Hamiltonian described above and count the number of states in this space at each value of angular momentum. Although the Hamiltonian is c-dependent, the state counting reproduces the character for a generic superconformal field theory (at arbitrary c), rather than any specific model.

I then try to modify the Hamiltonian to obtain states that agree with a theory of this type. In section 4.3.5, using related work by Feigin, Jimbo and Miwa [68], I show that one particular non-unitary superconformal minimal model ($SM(2,8)$) may be recovered by adding additional inter-cluster couplings to the states obtained for the generic SCFT in the previous section. In section 4.3.6 I attempt to employ the same procedure to obtain a unitary SCFT, the tricritical Ising model, and show that this is not possible, at least in the simplest scenario. I present results on an alternative approach involving a manual construction of the irreducible module from a different, larger overcomplete basis.

I offer some concluding remarks in section 4.4, along with some proposals for future directions this project should take.

4.1 Introduction to the fractional quantum Hall effect

The discovery and explanation of the quantum Hall effects is one of the major achievements of modern condensed matter physics. Reviewing any but the specific topics needed in the discussion of the results in section 4.3 would take us too far afield: the reader is referred to the reviews [120] and relevant chapters in [121], upon which the discussion below is largely based.

4.1.1 Integer and fractional effects

The quantum Hall effect (QHE) is a phenomenon observed in a cold (macroscopically coherent) two-dimensional electron gas with a current (along the $\mathbf{y}$ axis) applied perpendicular to a strong magnetic field (along the $\mathbf{z}$ axis). Such a system may be realized by, *e.g.*, the interface between two semiconducting materials with differing bandgaps and closely similar lattice constants. The QHE is experimentally characterized by a vanishing longitudinal resistivity $\rho_{xx} = 0$ and quantized Hall resistance $R_H = \rho_{xy} = \frac{1}{\nu}\frac{h}{e^2}$, where the filling fraction ν is a rational number. Note that R_H is independent of the sample geometry only in two dimensions, otherwise the rationality of ρ_{xy} would not be experimentally meaningful.

The physical mechanism behind this phenomenon is remarkable: the QHE states are topologically ordered states of matter which are incompressible in the bulk. Incompressibility means that

the chemical potential changes discontinuously with the density of the system — in other words, there is an energy gap for all excitations.

Integer quantum hall effect

Theoretical understanding of the QHE begins by neglecting the Couloumb interaction between electrons. The single-particle Hamiltonian is then simply

$$H_1 = \frac{1}{2m}\left(\mathbf{p} - \frac{e}{c}\mathbf{A}\right)^2 \tag{4.1}$$

and its eigenvalues are a set of degenerate, equally spaced bands known as Landau levels (as can be seen by solving H_1 in the Landau gauge $A_x = By, A_y = 0$). The Landau level spacing is given by the cyclotron energy $\hbar\omega_c = \frac{\hbar eB}{mc}$.

We consider first electrons in the disk geometry, occupying a droplet of area A near the origin. There are then $An_B = AB/\Phi_0$ single-particle states in each Landau level (in all of our discussion in this chapter we assume our particles are fully spin polarized), where the flux quantum $\Phi_0 = hc/e$. The filling fraction ν is then defined as n_e/n_B. The effect of adding an external E field to the single-particle Hamiltonian reproduces the classical Drude result

$$\sigma_{xy} = n_e\frac{ec}{B} = \nu\frac{e^2}{h} \tag{4.2}$$

If we exactly fill ν Landau levels, σ_{xy} appears to be quantized at integral values of $\frac{e^2}{h}$, and $\sigma_{xx} = 0$ since there are no nearby states for electrons to scatter into; however this does not explain the presence of plateaus of nonzero width where σ_{xy} remains constant. This arises by adding in the single-particle potential we have neglected so far: the effect of a lattice potential and disorder will be to broaden the Landau levels into bands of nonzero width, lifting the extensive degeneracy. However, in the presence of disorder all the states become localized except at the center of the band, and so the quantized conductance plateaus persist when the chemical potential of the electron gas lies in between the νth and $(\nu+1)$th Landau levels. The robustness of this effect is guaranteed by the fact that σ_{xy} is actually a topological invariant [122].

We have nonzero Hall conductance even when all accessible states are localized because the entire current is carried by edge states [123] which are chiral (due to the drift velocity imposed by the magnetic field) and hence dissipationless (no backscattering can occur if all the motion is in one direction). Hence localization and chiral edge states form the theoretical basis for the integer quantum Hall effect.

Fractional quantum hall effect

The preceding arguments should make the observed presence of plateaus at *fractional* values of ν all the more mysterious. We again look for the explanation in effects we have neglected: we are of course only justified in neglecting electron-electron interactions so long as the Landau level spacing $\hbar\omega_c$ is much larger than the Coulomb energy scale e^2/ℓ_B, where the magnetic length $\ell_B^2 = \frac{\hbar c}{eB}$. When a Landau level is only partially filled, these interactions may become important and a true many-body treatment of the problem is necessary. This is complicated by the fact that we don't have a nondegenerate soluble system to use as a starting point for perturbation theory, due to the extensive degeneracy of the many-particle basis states for a single Landau level.

Laughlin [50] made progress on the problem by introducing the philosophy of trial wavefunctions: instead of determining the Hamiltonian and writing its eigenfunctions, we instead guess an appropriate wavefunction which is compared with numerical exact diagonalization studies — the Hilbert space is finite-dimensional if we work in the lowest Landau level with a fixed number of electrons and flux quanta. Due to finite-size effects, this method is less than satisfactory: in particular, two states with different topological properties may not be numerically distinguishable for small numbers of particles. However, when the system is deep within a quantum Hall phase, correlation lengths should be small and we should be able to obtain a fair description from few particles. More support for this procedure was provided by Haldane's pseudopotential argument [124], which I summarize in section 4.1.3 after first establishing a framework for discussing trial wavefunctions.

4.1.2 Lowest Landau level polynomials

If we adopt the symmetric gauge $\mathbf{A} = \frac{1}{2}\mathbf{B} \times \mathbf{r}$ and use complex coordinates $z = r_1 + ir_2$ and $\overline{z} = r_1 - ir_2$, the eigenfunctions of the single-particle Hamiltonian in the lowest Landau level (LLL) are given by

$$u_m(z, \overline{z}) \propto z^m \exp -\frac{|z|^2}{4\ell_B^2}, \tag{4.3}$$

where the magnetic length $\ell_B^2 = \frac{\hbar c}{eB}$. These are angular momentum eigenstates, localized in a ring with radius of order $\sqrt{2m}\ell_B$ around the origin. In the remainder of our discussion, for convenience I choose units so that $\ell_B = 1$. Multi-particle states can be expressed in a basis of products of the single-particle states, so wavefunctions in the lowest Landau level correspond to analytic polynomials multiplied by the Gaussian factor [125].

Frequently it becomes to consider quantum Hall states on compact manifolds, so that edge effects are irrelevant. Following [124], we can consider wavefunctions on a sphere of radius R, whose surface is penetrated by N_ϕ quanta of magnetic flux (which we can imagine as coming from a monopole in the interior). In the LLL each particle has angular momentum $N_\phi/2$. One may write wavefunctions on the sphere in terms of homogenous (or spinor) coordinates u, v defined in terms of the spherical coordinates θ, ϕ as $u = e^{i\phi/2}\cos\theta/2$, $v = e^{-i\phi/2}\sin\theta/2$. We may eliminate the redundancy involved in this description by stereographic projection [66], choosing $z = 2Rv/u$. The LLL single particle basis states are then

$$\varphi_m(z) \propto \frac{z^m}{\left(1 + \frac{|z|^2}{4R^2}\right)^{1+N_\phi/2}}, \quad \text{for } m \leq N_\phi, \tag{4.4}$$

with the latter constraint added to make the inner product $\langle\varphi_i|\varphi_j\rangle = \int dz\, dz^*\, \varphi_i^*(z)\varphi_j(z)$ well-defined. We see that each particle behaves as a spin, with $L_z = N_\phi/2 - m$.

We can go from the sphere geometry to the disk by taking the radius of the sphere to be very large while keeping the flux density constant. With $R \to \infty$, $N_\phi/R^2 \to$ const., we recover the planar single particle functions (4.3). This procedure also defines a shift S via

$$N_\phi = \nu^{-1}N - S. \tag{4.5}$$

The discussion to follow will center on states on the plane and sphere, so we may specify a many-body state as an analytic polynomial in the particle coordinates. Statistics dictate that this function is (anti)symmetric for (fermions) bosons. The fermionic case may be reduced to the bosonic case by extracting a Vandermonde determinant:

$$\psi_F(z_1, \ldots, z_k) = \prod_{i<j}(z_i - z_j)\psi_B(z_1, \ldots, z_k). \tag{4.6}$$

where ψ_F, ψ_B are any (anti)symmetric homogenous polynomials, so we need only consider the bosonic case from here on. The filling fractions in the two cases obey

$$\nu_F = \frac{\nu_B}{\nu_B + 1} \tag{4.7}$$

Note that quantum Hall states of bosons may be directly experimentally realizable by rapidly rotating condensates of cold atomic gases [126].

4.1.3 Laughlin states and pseudopotentials

Pseudopotential formalism

The many-body Hamiltonian projected to the lowest Landau level consists only of the interaction term, which must be translationally and rotationally invariant. The many-particle wavefunction may be decomposed into two-particle wavefunctions as

$$\Psi(r_1, \ldots, r_N) = \sum_i \psi_i(r_1, r_2)\widetilde{\Psi}_i(r_3, \ldots, r_N). \tag{4.8}$$

This two-body function may, in turn, be decomposed in terms of functions depending only on the center of mass coordinate $(r_1 + r_2)/2$ and the relative separation $r_1 - r_2$. Since the interaction is translationally invariant it will only act on the relative piece.

An appropriate basis for the relative wavefunctions is provided, in the lowest Landau level, by the powers

$$|L\rangle = (z_1 - z_2)^L, \tag{4.9}$$

where L is the relative angular momentum between the two particles in the plane. In order to reproduce the proper statistics when particles are interchanged, L must be odd (even) for fermions (bosons). If the interaction is rotationally invariant, it will not couple states of different angular momentum. Given this basis, we can decompose any interaction $V(r_1 - r_2)$ as

$$H_{12} = \sum_L |L\rangle\langle L|V(r_1 - r_2)|L\rangle\langle L| \tag{4.10}$$

so that the full Hamiltonian may be written as

$$H = \sum_{i<j} \sum_L V_L P_{ij}^L \tag{4.11}$$

where $P_{ij}^L = |L_{ij}\rangle\langle L_{ij}|$ projects onto the state where particles i and j have relative angular momentum L, and V_L are the psedupotential coefficients obtained from the above inner product. Note that may potentials may give rise to the same set of V_L, so these coefficients may be thought of as capturing those aspects of the interaction that are relevant to lowest Landau level physics. These are the eigenenergies of the two-particle system; for more particles relative angular momentum is no longer a good quantum number since, e.g., the operators L_{12} and L_{13} do not commute.

Laughlin states

It would appear that these manipulations have brought us no closer to being able to deal with the presence of two-body interactions. However, as Haldane pointed out [124], if the potential takes a "hard-core" form with $V_L = 0$ for $L \geq M$, $V_L > 0$ for $L < M$, the highest-density (lowest angular momentum at fixed particle number) eigenstate may be found exactly: since

$$P_{ij}^L (z_i - z_j)^M = 0 \tag{4.12}$$

for $L \geq M$, the Hamiltonian vanishes when acting on

$$\Psi_M(z_1, \ldots, z_N) = \prod_{i<j} (z_i - z_j)^M e^{-\frac{1}{4}\sum_i |z_i|^2} \tag{4.13}$$

where M odd (even) yields Fermi (Bose) statistics. This is the original trial wavefunction written down by Laughlin [50], which describes a quantum Hall system at filling fraction $\nu = 1/M$. (There is an important caveat that the particles must be distributed in a macroscopically uniform manner in order for ν to be well defined, which Laughlin argued by appealing to a Coulomb plasma analogy). Because it is the unique state of minimal total angular momentum (equal to $\frac{1}{2}MN(N-1)$, which can be seen from power counting in (4.13)), it is incompressible: shrinking the droplet would require a reduction of angular momentum.

The pure Couloumb interaction, of course, is not of this hard-core form, but its pseudopotentials V_L decrease rapidly with increasing L, so the hard-core Hamiltonian is a sensible approximation. Ultimately, though, its validity is established through numerical diagonalization, which displays remarkably high overlap with the Laughlin state — in fact, it is not even known analytically whether the Haldane Hamiltonian remains gapped in the thermodynamic limit.

The Laughlin state is well-established at the time of writing as describing the essential physics of the $\nu = 1/M$ states. However, other plateaus are observed at values of ν which are not of this form. The explanation of these states is less unambiguous, but one well-studied proposal is provided by the hierarchy construction [127]. Here one begins with a Laughlin state and changes ν by creating an extensive number of quasihole or quasielectron excitations (to be described below), which then may organize themselves into a Laughlin state. This state may have excitations in turn, which can organize themselves into still another state, and so on. The reader should be aware that this is only one of several competing proposals; I do not summarize them here since I will not be concerned with these states in what follows.

Excitations of the Laughlin states

The Laughlin state is the highest-density zero-energy eigenstate of the hard-core Haldane Hamiltonian. However, the argument also goes through if we multiply the Laughlin state by any symmetric polynomial. We may organize these in the form of a quasihole factor

$$\Psi^{+}_{M;w}(z_1, \ldots, z_N) \propto \prod_i (z_i - w) \prod_{i<j} (z_i - z_j)^M e^{-\frac{1}{4}\sum_i |z_i|^2}. \tag{4.14}$$

(where I have neglected a w-dependent normalization factor which is a matter of convenience). This describes a quasihole since the wavefunction now vanishes as any particle approaches w. Because the electron liquid is incompressible, the particles shift radially outward from this position and the particle density away from w is unchanged, so the quasihole appears to be a localized region of *fractional* charge e/M (if we had M quasiholes in the same location, they would appear to be another particle — this may be verified in the Coulomb plasma analogy). Because of this fractional charge, the quasiholes have fractional statistics [128]: we accumulate an Aharonov-Bohm phase of $\theta = \pi/M$ when two quasiholes are adiabatically exchanged.

The presence of fractionally charged excitations should not be cause for confusion: again due to incompressiblity, we may view the excess charge $-e/M$ as accumulating at the edge of the droplet, since the droplet as a whole is now larger. In this way quasiholes may also be used to describe the gapless edge excitations of the droplet, by taking w far outside the droplet (whose radius is $\sqrt{2M(N-1)}$, as can be inferred from the fact that a single-particle state of angular momentum (degree) m is localized at radius $\sqrt{2m}$) [129].

One can also consider quasiparticle excitations, which would schematically be created by factors of $\prod_i (z_i - w)^{-1}$. However this factor is non-analytic and hence takes us out of the LLL. Instead [50] we act on the Jastrow factor with a differential operator:

$$\Psi^{-}_{M;\overline{w}}(z_1, \ldots, z_N) \propto \left[\prod_i \left(2\frac{\partial}{\partial z_i} - \overline{w} \right) \prod_{i<j} (z_i - z_j)^M \right] e^{-\frac{1}{4}\sum_i |z_i|^2}. \tag{4.15}$$

Note that this wavefunction is not analytic in the quasiparticle coordinate, but it does not need to be: the ws are only static parameters describing the locations of the excitations. As we saw in section 4.1.2, the wavefunction only has to be analytic in the actual particle coordinates. Because of the difficulty in dealing with differential operators, I will only concern myself with quasihole excitations in what follows.

4.1.4 FQHE states, Chern-Simons theory and conformal blocks

My presentation of the FQHE so far should appear intellectually unsatisfying to the reader: surely there is more physical content to the subject than simply selecting symmetric polynomials at random and numerically testing the resulting wavefunction! In fact, the fractional quantum Hall effect provides a realization of very deep physical concepts, as should be clear from the appearance of fractional statstics in the preceding section; I have emphasized the trial wavefunction viewpoint because it forms the foundation for my arguments in the following sections.

Girvin and MacDonald [130] (see also [131]) first proposed that the Laughlin state manifested long-range order, but with respect to a nonlocal order parameter. A similar argument was given independently by Read [132], who proposed the following order parameter operator:

$$\Psi^\dagger(z) = \psi^\dagger(z) \prod_i (z_i - z)^M e^{-|z|^2/4} \tag{4.16}$$

where $\psi^\dagger(z)$ creates a particle at z. The Laughlin state at filling faction $1/M$ may be obtained as a Bose condensate of $\Psi^\dagger$ (the combination of one particle and M quasiholes, or flux quanta, has bosonic statistics). The Landau-Ginzburg effective field theory of the order parameter $\langle \Psi^\dagger \rangle$ contains a Chern-Simons (CS) term, which should not be surprising after the discussion of 4.2.4: there we saw that anyonic statistics necessarily required the presence of such a term.

The full implications of this were developed in [119]. As we will see in section 4.2.4, the wavefunctions of a CS theory are (chiral) conformal blocks of an associated WZW conformal field theory. Moore and Read proposed to use these blocks as trial wavefunctions for FQHE states: particles and quasiparticles would correspond to insertions of various operators of the CFT, with their statistics being determined by the corresponding braiding operator braiding properties; this was put on a more rigorous footing in [133].

It was first realized in [119] that, through this correspondence, one could also construct FQHE states whose excitations possess *non-Abelian* statistics. As an example, a trial wavefunction (the Moore-Read or Pfaffian state) was given in that paper which was constructed from blocks of the $M(3,4)$ conformal field theory (which corresponds to the $d = 2$ Ising model). There is currently

some experimental confirmation that this wavefunction correctly describes the state at $\nu = 5/2$. Other physically relevant non-Abelian states have been proposed (*e.g.*, [62]). Non-Abelian FQHE states currently constitute the most practical scheme for implementing topological quantum computing [134].

4.2 Summary of conformal field theory

In this section I introduce various aspects of conformal field theory that will be useful in the following discussion. It is beyond the scope of this thesis to give a completely self-contained review; the discussion that follows will be extremely condensed. Fortunately many introductions to the subject of varying length and rigor exist; I direct the reader to [135, 136]. Many of the results of this section were first obtained by Belavin, Polyakov and Zamolodchikov [137].

4.2.1 Conformal invariance

The group of conformal transformations (colloquially speaking, those that preserve angles) in d dimensions is isomorphic to the (non-compact) group $SO_+(d+1,1)$. Since this is larger than the group of scale, rotation and translation transformations, it places more constraints on the correlation functions of observables of our theory. Recall (*e.g.*, section 3.3.3) that scale invariance implies

$$\langle \phi_1(\mathbf{r}_1)\phi_2(\mathbf{r}_2)\cdots\rangle = \prod_i b^{x_i}\langle \phi_1(b\mathbf{r}_1)\phi_2(b\mathbf{r}_2)\cdots\rangle. \tag{4.17}$$

We will define a quasi-primary field (primary fields will be defined below) as one that transforms as

$$\phi(\mathbf{r}) \mapsto \phi'(\mathbf{r}') = \left|\frac{\partial \mathbf{r}'}{\partial \mathbf{r}}\right|^{-x/d} \phi(\mathbf{r}) \tag{4.18}$$

under a global conformal transformation $\mathbf{r} \mapsto \mathbf{r}'$, where $|\partial\mathbf{r}'/\partial\mathbf{r}|$ is the Jacobian of that transformation. This fixes the form of the two-point function of such fields to

$$\begin{aligned}\langle \phi_i(\mathbf{r}_1)\phi_j(\mathbf{r}_2)\rangle &= \frac{c_{ij}}{|\mathbf{r}_1-\mathbf{r}_2|^{2x_i}} \text{ if } x_i = x_j,\\ &= 0 \text{ otherwise.}\end{aligned} \tag{4.19}$$

By suitable orthonormalization in the space of fields, we may take the constants $c_{ij} = \delta_{ij}$. The form of the three-point function is also fixed as

$$\langle \phi_i(\mathbf{r}_1)\phi_j(\mathbf{r}_2)\phi_k(\mathbf{r}_2)\rangle = \frac{c_{ijk}}{\mathbf{r}_{12}^{x_1+x_2-x_3}\mathbf{r}_{23}^{x_2+x_3-x_1}\mathbf{r}_{31}^{x_3+x_1-x_2}} \tag{4.20}$$

for undetermined constants c_{ijk}, where $\mathbf{r}_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$. Higher-order correlations will be determined up to some functions of the conformally-invariant cross-ratios $\mathbf{r}_{ij}\mathbf{r}_{kl}/\mathbf{r}_{ik}\mathbf{r}_{jl}$.

Crucially, in two (and only two) dimensions the group of *local* conformal transformations, which need not be well-defined everywhere, is infinite-dimensional. If we adopt complex coordinates $z = r_1 + ir_2$ and $\overline{z} = r_1 - ir_2$, the global conformal transformations in $d = 2$ are the Möbius transformations (isomorphic to $SL(2,\mathbb{C})/\mathbb{Z}_2$) and any analytic mapping $z \mapsto w(z), \overline{z} \mapsto \overline{w}(\overline{z})$ is locally a conformal transformation. Equation (4.18) becomes

$$\phi(z,\overline{z}) \mapsto \phi'(w,\overline{w}) = \left(\frac{dw}{dz}\right)^{-h}\left(\frac{d\overline{w}}{d\overline{z}}\right)^{-\overline{h}}\phi(z,\overline{z}) \tag{4.21}$$

for two real numbers $h, \overline{h}$ which are *not* complex conjugates of each other. The scaling dimension $x = h + \overline{h}$, while the difference $h - \overline{h}$ describes the (conformal) spin of the field. If equation (4.21) holds for any *local* conformal transformation, in addition to the global transformations, we refer to ϕ as a primary field. These are in some sense the most local objects in a conformal field theory, and their behavior essentially fixes its structure (in a way to be made more specific in section 4.2.4).

Conformal invariance also fixes the form of the operator product expansion (OPE). Recall that in any correlation function of a scale invariant theory we have

$$\langle \phi_i(\mathbf{r}_1)\phi_j(\mathbf{r}_2)\cdots\rangle \sim \sum_k C_{ij}^k(\mathbf{r}_{12})\langle \mathcal{O}_k(\mathbf{r}_2)\cdots\rangle \tag{4.22}$$

as $\mathbf{r}_1 \to \mathbf{r}_2$, for some functions C_{ij}^k and new operators $\mathcal{O}_k$. This merely expresses the fact that two fields acting close together will appear to act as a single operator when viewed from far away; an example was given by the mass insertion field ϕ^2 in percolation (section 3.3.2). Conformal

invariance in two dimensions makes this relation much more quantitative: we have

$$\langle \phi_i(z,\overline{z})\phi_j(w,\overline{w})\cdots\rangle \sim \sum_k \sum_{\{p\},\{\overline{p}\}} c_{ijk}\beta_{ijk}^{\{p\}}\beta_{ijk}^{\{\overline{p}\}} \times (z-w)^{h_k+\sum\{p\}-h_i-h_j}(\overline{z}-\overline{w})^{\overline{h}_k+\sum\{\overline{p}\}-\overline{h}_i-\overline{h}_j}\langle \phi_k^{\{p\},\{\overline{p}\}}(w,\overline{w})\cdots\rangle \quad (4.23)$$

where the structure constants c_{ijk} are the same as those appearing in the three-point function (4.20), and the constants β are uniquely defined by fixing h_i, h_j, h_k. The additional sets of integers $\{p\}$ and $\{\overline{p}\}$ signify that the fields ϕ_k appearing on the right-hand side are not all primary. The chief importance of (4.23) is that the OPE now has a finite radius of convergence (up to the insertion of the next nearest operator).

The stress tensor and Ward identities

Because we wish to study the effects of local coordinate transforms on our theory, we are necessarily led to consider the properties of the stress-energy tensor $T_{\mu\nu}$. Recall that this object describes the response of the action of the theory to an infinitesimal change in metric, and that translation, rotation and scale invariance imply that it is a conserved Noether current and that it is a symmetric traceless tensor (we are considering what is referred to by some sources as the "improved" stress-energy tensor). This leaves two independent components in $d = 2$; if we parameterize these as $T = T_{11} + iT_{12}$ and $\overline{T} = T_{11} - iT_{12}$ the conservation laws specify that they are holomorphic and antiholomorphic, respectively.

We now consider the effect of an infinitesimal conformal transformation $z \mapsto w = z + \varepsilon(z)$ (and similarly for $\overline{z}$). The definition of a primary field (4.21) implies

$$\delta_{\varepsilon,\overline{\varepsilon}}\phi(w,\overline{w}) = h\phi(z,\overline{z})\partial\varepsilon(z) + \varepsilon(z)\partial\phi(z,\overline{z}) + \text{CC}, \quad (4.24)$$

where CC denotes the complex conjugate of the rest of the right-hand side. We can also view this

transformation as being generated by the stress tensor (as a consequence of its definition), via

$$\delta_{\varepsilon,\bar{\varepsilon}}\phi(w,\overline{w}) = \frac{1}{2\pi i}\oint_w [dz\, T(z)\varepsilon(z), \phi(w,\overline{w})] + \text{CC}, \tag{4.25}$$

where the contour in the z plane is a small circle around w. Equating these two expressions gives the OPE for any conformal field and the stress tensor,

$$T(z)\phi(w,\overline{w}) \sim \frac{h\phi(w,\overline{w})}{(z-w)^2} + \frac{\partial\phi(w,\overline{w})}{z-w} \tag{4.26}$$

and analogously for $\overline{T}(\overline{z})\phi(w,\overline{w})$. This holds at the level of the integrand of (4.25), since the equality must hold for arbitrary transformations. Repeating this argument for an infinitesimal conformal transformation of a bounded region containing a number of primary fields, we obtain

$$\begin{aligned}\langle T(z)\phi_1(w_1,\overline{w}_1)\cdots\phi_n(w_n,\overline{w}_n)\rangle \\ \sim \sum_{i=1}^{n}\left[\frac{h_i}{(z-w_i)^2} + \frac{\partial_{w_i}}{z-w_i}\right]\langle\phi_1(w_1,\overline{w}_1)\cdots\phi_n(w_n,\overline{w}_n)\rangle.\end{aligned} \tag{4.27}$$

This is the conformal Ward identity, which may be taken as axiomatic when dealing with CFTs that are not defined by a classical action.

We now consider the special case where the transformation in question is a global transformation. This means $\varepsilon(z)$ will be a polynomial of degree less than or equal to two in z, and we can take the domain of the transformation to be the entire complex plane and send $z \to \infty$ in the integrated form of (4.27). We require all correlation functions to be invariant under these transformations, so in order for the right-hand side of (4.27) to vanish we need $T(z)$ to fall off at least as fast as z^{-4} as $z \to \infty$. Note that this is a stronger constraint than the z^{-2} behavior dictated by its scaling dimension. This means that the stress tensor is not a primary field; it must instead have the following OPE with itself:

$$T(z)T(w) \sim \frac{\frac{1}{2}c}{(z-w)^4} + \frac{2T(w)}{(z-w)^2} + \frac{\partial T(w)}{z-w}, \tag{4.28}$$

and likewise for $\overline{T}$, with $T(z)\overline{T}(\overline{w})$ being nonsingular. This is a consequence of the above behavior

and the fact that T has weights $h = 2, \overline{h} = 0$ as a consequence of its scaling and rotational properties. Here the constant c is the central charge of the theory, which is determined by the actual physical content of the theory rather than considerations of conformal invariance alone.

We then obtain the following consequences of the Ward identity (4.27) for the case of global conformal transformations:

$$\sum_{i=1}^{n} \partial_{z_i} \langle \phi_1(z_1) \cdots \phi_n(z_n) \rangle = 0 \tag{4.29}$$

$$\sum_{i=1}^{n} \left(z_i \partial_{z_i} + h_i \right) \langle \phi_1(z_1) \cdots \phi_n(z_n) \rangle = 0 \tag{4.30}$$

$$\sum_{i=1}^{n} \left(z_i^2 \partial_{z_i} + 2 z_i h_i \right) \langle \phi_1(z_1) \cdots \phi_n(z_n) \rangle = 0 \tag{4.31}$$

Note that these also imply the forms of the two- and three-point functions found in (4.19), (4.20).

4.2.2 Radial quantization

We quantize a two-dimensional conformally invariant field theory by capitalizing on its scale invariance. The exponential map sends the complex plane to an infinite strip of width 2π with periodic boundary conditions. Translations along the infinite (time) direction in the former correspond to scalings in the latter, so the Hamiltonian for our theory will be the generator of scale transformations. We must then *radially* order the action of various operators.

We may expand $T(z), \overline{T}(\overline{z})$ in (operator-valued) Fourier modes

$$T(z) = \sum_{n=-\infty}^{\infty} L_n z^{-n-2}; \qquad L_n = \frac{1}{2\pi i} \oint_0 dz\, T(z) z^{n+1}, \tag{4.32}$$

and analogously for $\overline{T}$ (in what follows I will frequently neglect the antiholomorphic half of the equations, since the discussion exactly parallels the holomorphic case.) The modes L_n are the generators of infinitesimal conformal transformations $\varepsilon(z) = z^{n+1}$. Their algebra may be determined by considering the application of two such transformations in turn, under radial ordering, using the

Ward identity (4.27) and the OPE (4.28). We obtain

$$[L_m, L_n] = (m-n)L_{m+n} + \frac{c}{12}m(m^2-1)\delta_{m+n,0} \tag{4.33}$$

and analogously for $[\overline{L}_m, \overline{L}_n]$, with $[L_m, \overline{L}_n] = 0$. The relations (4.33) define the Virasoro algebra [138]. Because $T(z)$ is a hermitian field, we also have

$$L_n^\dagger = L_{-n}. \tag{4.34}$$

Inspection of (4.33) shows that L_0, L_1, L_{-1} form a closed subalgebra, which generates the global conformal transformations. The full group of local conformal transformations is generated by the direct product of two copies of the Virasoro algebra (the Ls and the $\overline{L}$s), so it will be productive to express the Hilbert space of our theory in terms of irreducible representation of this algebra. Since the Ls and $\overline{L}$s commute, all such representations will be tensor products of two representations of the Virasoro algebra, so in what follows we may restrict our attention to just the holomorphic half.

Representations of the Virasoro algebra

Scale transformations in the plane are generated by $L_0 + \overline{L}_0$, which therefore plays the role of Hamiltonian (generator of infinitesimal time translations) for our theory. Because $[L_0, L_n] = -nL_n$, we may regard the half of the Virasoro algebra with $n < 0$ as raising operators and the half with $n > 0$ as lowering operators. We label eigenstates $|h\rangle$ of L_0 by their eigenvalues, or weights, h. These states are lowest-weight states if $L_n|h\rangle = 0$ for all $n > 0$: the entire Hilbert space of the theory is generated by acting on the lowest-weight states with combinations of the raising operators L_{-n} (analogously to the treatment of representations of $su(2)$ in ordinary quantum mechanics).

This produces a conformal tower of descendant states, graded by the eigenvalues of L_0:

$$\begin{aligned} |h\rangle \\ L_{-1}|h\rangle \\ L_{-2}|h\rangle, L_{-1}^2|\phi\rangle \\ L_{-3}|h\rangle, L_{-2}L_{-1}|\phi\rangle, L_{-1}^3|\phi\rangle \\ \dots \end{aligned} \tag{4.35}$$

Note that because of the commutation relation (4.33) we may always order the raising operators so that those with lower indices act first.

Conformal field theory manifests an operator-state correspondence: each of the lowest-weight states $|h\rangle$ corresponds to the insertion of a particular primary field ϕ with weight h at the origin in the plane (so that it is inserted to the right of all other fields in any radially-ordered correlation function). The vacuum $|0\rangle$ of our theory (which we will assume is nondegenerate) should be invariant under the global conformal transformations, hence $|0\rangle$ should be annihilated by L_1, L_0 and L_{-1} — this ensures the invariance of all correlation functions under these transformations.

It may happen that a state $|\chi\rangle$ at some level of the module is, itself, a highest weight state; i.e. that L_n acting on that state for all $n > 0$ vanishes. If this happens, $|\chi\rangle$ is a null (or singular) vector, and so are all its descendants under the action of the Virasoro generators. Descendants of a null vector have zero inner product with all other states in the module, and hence in order to obtain a unitary theory (with positive definite inner product between all states) all null states must be removed, yielding an irreducbile representation of the Virasoro module (*i.e.*, one which cannot be decomposed into submodules).

For example, we have a null state at level two if

$$L_n \left(L_{-2} - \frac{1}{\alpha} L_{-1}^2 \right) |h\rangle = 0, \qquad n > 0 \tag{4.36}$$

for some value of α. We can examine the specific cases $n = 1$, $n = 2$ and deduce that this requires

$$h = \frac{3\alpha - 2}{4}; \qquad c = \frac{(3\alpha - 2)(3 - 2\alpha)}{\alpha}. \tag{4.37}$$

This sort of argument was extended by Kac [139], who examined the Gram matrix of the Verma module states at arbitrary level and found that it vanished (*i.e.*, that a null state existed) at level rs if

$$h = h_{r,s}(\alpha) = \frac{(r\alpha - s)^2 - (\alpha - 1)^2}{4\alpha}, \tag{4.38}$$

with $c = c(\alpha)$ as given in (4.37). This result was proved by Feigin and Fuchs [140].

The presence of singular vectors gives us useful differential equations for correlation functions of primary operators. As a consequence of the Ward identity (4.27) and the primary field OPE (4.26), for any set of primary fields we have

$$\langle L_{-m}\phi(z)\phi_1(w_1)\cdots\phi_n(w_n)\rangle = \mathcal{L}_{-m}\langle\phi(z)\phi_1(w_1)\cdots\phi_n(w_n)\rangle \tag{4.39}$$

where

$$\mathcal{L}_{-m} = \sum_{i=1}^{n} \frac{(m-1)h_i}{(z-w_i)^m} - \frac{\partial_{w_i}}{(z-w_i)^{m-1}}. \tag{4.40}$$

If ϕ is a singular operator, we know any correlation function involving it must vanish, which gives us a differential equation for correlation functions involving the primary field from which it was derived (its "ancestor"). For the level two example (4.36), this implies

$$\left[\mathcal{L}_{-2} - \frac{1}{\alpha}\mathcal{L}_{-1}^2\right]\langle\phi(z)\phi_1(w_1)\cdots\phi_n(w_n)\rangle = 0. \tag{4.41}$$

When these relations are applied to three-point functions, we obtain constraints on the coefficients c_{ijk} defined in (4.20). This in turn limits the possible operators appearing on the right-hand side of the OPE of the ancestor operator with other primary fields.

4.2.3 Minimal models

If α is rational and equal to p/p' (with p, p' relatively prime), the Kac formula (4.38) becomes

$$h_{r,s} = \frac{(rp - sp')^2 - (p - p')^2}{4pp'} \tag{4.42}$$

for a theory of central charge

$$c = 1 - 6\frac{(p - p')^2}{pp'} \tag{4.43}$$

which has the symmetry $h_{r,s} = h_{p'-r,p-s}$. Note without loss of generality we may take $p > p'$. As a consequence of the singular vector constraints on the OPEs mentioned above, it turns out that the OPEs of operators in the region $1 \leq r < p', 1 \leq s < p$ only involve other operators from the same set. This region is called the Kac table, and it follows that, at rational central charge, theories can be constructed with only a finite number of primary fields. A theory consisting of these fields with central charge and weights given by (4.43), (4.42) is referred to as a (Virasoro) minimal model, or more specifically as $M(p, p')$.

It was shown by Friedan, Qiu and Shenker [141] that, of the Virasoro minimal models, the only unitary ones have $p \geq 3$ and $|p - p'| = 1$ (this may be seen nonrigorously by requiring the two-point function (4.19) to vanish for large separations, which requires all $h_{r,s}$ to be positive). This is the series of Virasoro unitary minimal models, and they completely classify the unitary CFTs with $0 < c < 1$. No unitary theories exist with $c < 0$; the set of unitary theories at $c = 1$ has also been classified (see [142]). The classification of unitary theories with $c > 1$ is still an open problem, to which the formalism of rational conformal field theory (section 4.2.4) provides a partial answer.

Note that we obtained these conclusions as a consequence of the Virasoro algebra. If our theory may be classified according to representations of some larger algebra containing (4.33) as a subalgebra, we will obtain different OPE constraints and a different set of possible minimal models; an example will be given below.

Character formulae

Having constructed representations of the Virasoro algebra in terms of Verma modules, a natural thing to do is count the dimension of the space of states at a given level. This information is encoded in a generating function known as the character of the module. For a generic Verma module (one without singular vectors) it follows from (4.35),

$$\chi(q) = \sum_{n=0}^{\infty} d(n) q^n = \prod_{m=1}^{\infty} \frac{1}{1-q^m} \equiv \frac{1}{(q)_\infty}. \tag{4.44}$$

The fact that the coefficient of q^n on the right-side of (4.44) count the number of integer partitions of n (see section 4.3.1) can be seen from writing each factor $1/(1-q^m)$ as an infinite geometric series in q^m and expanding out all terms.

The presence of singular vectors significantly complicates this counting exercise. The process of constructing the irreducible representations of the Virasoro algebra for the minimal models is quite technically involved: as it turns out, the modules generated by a singular vector contain singular vectors of their own in turn, so all of these must be removed by an inclusion-exclusion process [140]. I simply quote the result obtained in [143]: the character for the irreducible Verma module of the highest-weight state $|h_{r,s}\rangle$ of the $M(p,p')$ Virasoro minimal model (with $h_{r,s}$ defined in (4.42)) is

$$\chi^{p,p'}_{r,s}(q) = K^{p,p'}_{r,s}(q) - K^{p,p'}_{r,-s}(q). \tag{4.45}$$

where I have defined

$$K^{p,p'}_{r,s}(q) = \frac{q^{-1/24}}{(q)_\infty} \sum_{n=-\infty}^{\infty} q^{(2pp'n+pr-p's)^2/(4pp')}. \tag{4.46}$$

Fermionic characters

The result (4.45) applies to all primary fields of all Virasoro minimal models, but that very fact may also be seen as a weakness: it tells us very little about the actual physics behind each model. An improvement is offered in the form of fermionic characters, which are expressed in terms of a sum of terms which are all manifestly positive.

4.2.4 TQFT, RCFT and all that

As I noted in section 4.2.3, the most useful conformal field theories are those with a finite number of primary fields; *i.e.*, a finite number of irreducible highest-weight representations under some extended algebra including the Virasoro algebra. One powerful way to classify these theories was provided by Moore and Sieberg [144, 145] based on results of Witten [146], by appealing to topological field theories living in $2 + 1$ dimensions.

The application to fractional quantum Hall states was developed, essentially completely, in [119], and will be discussed later in section 4.1.4. These results are summarized in the recent review [134], which I will follow closely below.

Anyons and the braid group

The concept of particle statistics, familiar from elementary quantum mechanics, deals with the effect on a wavefunction of several identical particles when the positions of two of those particles are exchanged. In a space-time formalism, a continuous (in fact, this needs to be adiabatic) motion of two particles that exchanges their positions has the effect of topologically braiding their world-lines. This means that, under exchange, the wavefunction must transform as some (projective) representation of the braid group. The braid group on N strands (particle worldlines) is generated by exchanges of adjacent particles and is non-Abelian.

In one spatial dimension, particle statistics is not a well-defined concept since the exchange of identical particles is not well-defined either. In three and higher dimensions, if we perform two exchanges of the same chirality we may always deform the resulting braid to the trivial one. This has the effect of imposing an additional relation on the generators of the braid group, causing it to reduce to the permutation group (on N objects). There are only two possible one-dimensional representations of the permutation group, and these yield the familiar Bose and Fermi statistics. Higher-dimensional representations (parastatistics) are possible, but this may be reduced to bosons and fermions carrying extra quantum numbers.

The situation is much richer in two dimensions. Here we cannot make the deformation described, which has important consequences: while the permutation group is of finite order for finite

N, the braid group is of infinite order for any number of particles. It was first realized in [147] that the allowed one-dimensional representations of such a group could yield an arbitrary phase under exchange, a situation referred to as anyonic statistics.

Furthermore, higher-dimensional representations of the braid group cannot be related to the one-dimensional representations. This is the case of nonabelian statistics [144, 146]. Here there is a finite-dimensional space of wavefunctions for fixed particle positions, and a braiding operation has the effect of a linear transformation on that space. Perhaps the most amazing aspect of the fractional quantum Hall effect is that this seemingly abstract representation-theoretic concept is apparently implemented in nature.

Chern-Simons theory and TQFT

It then becomes natural to ask how anyonic statistics may be implemented in a field theory. We can do this by making use of the Aharonov-Bohm effect: if we couple ordinary particles to a gauge field, they will pick up an additional phase as they are exchanged adiabatically. This means that the gauge flux must be rigidly coupled to the particles and follow their positions when they are exchanged, which means the usual Maxwell term is unsuitable. Instead we must introduce a Chern-Simons (CS) term [146]: given a Lagrangian $\mathcal{L}_0$ with a conserved current j^μ, we add the terms

$$\mathcal{L} = \mathcal{L}_0 + \frac{k}{4\pi}\epsilon^{\mu\nu\rho}a_\mu\partial_\nu a_\rho + a_\mu j^\mu \tag{4.47}$$

where a is (for now) a $U(1)$ gauge field. The equation of motion for a is $\frac{k}{2\pi}\epsilon^{\mu\nu\rho}\partial_\nu a_\rho + j^\mu = 0$; the density of a collection of stationary particles is j^0, so this means that each particle has a flux of $2\pi/k$ attached to it. Under adiabatic exchange, this leads to an Aharonov-Bohm phase of $\theta = \pi/k$.

Adding the CS term has no other dynamical consequences for the theory; two components of a may be fixed in $2+1$ dimensions by a gauge choice and the final component is fixed by the equation of motion. The field a has no dynamics and no local degrees of freedom. This means that the pure CS theory is an example of a topological field theory, where all observables are diffeomorphism-invariant (*i.e.*, do not depend on the choice of coordinates). This can be seen from the fact that indices are contracted in the CS term without use of the metric $g_{\mu\nu}$ and the combination $d^3\mathbf{r}\,\epsilon^{\mu\nu\rho}$

is a tensor density of weight zero (so no factors of $\sqrt{-g}$ enter).

Rational CFTs from CS theory

We may obtain a two-dimensional conformal field theory (in fact, all known two-dimensional conformal field theories) from a topological field theory as follows (following Witten [146]). We start with a pure Chern-Simons (CS) action on a three-dimensional manifold M:

$$S = \frac{k}{4\pi} \int_M d^3r \, \epsilon^{\mu\nu\rho} \left(a^\alpha_\mu \partial_\nu a^\alpha_\rho + \frac{2}{3} f_{\alpha\beta\gamma} a^\alpha_\mu a^\beta_\nu a^\gamma_\rho \right) \tag{4.48}$$

where the gauge field a takes values in the Lie algebra of some (not necessarily Abelian) group G, with structure constants $f_{\alpha\beta\gamma}$. This is a topological field theory, since it may be constructed without reference to a metric tensor (the density $d^3r \, \epsilon^{\mu\nu\rho}$ is a diffeomorphism invariant).

The field a behaves as $a_\mu \mapsto g(a_\mu + \partial_\mu)g^{-1}$ under gauge transformations, where g is any function defined on the manifold M taking values in G. As with any pure gauge theory, observables will correspond to Wilson loops Γ:

$$W_R(\Gamma) = \mathrm{Tr}\, \mathbb{P} \exp i \oint_\Gamma T^\alpha_R a^\alpha_\mu dx^\mu \tag{4.49}$$

where $\mathbb{P}$ denotes a path-ordering operation and T^α are the generators of the algebra in some representation R. Correlation functions of several Wilson loops

$$\left\langle \prod_i W_{R_i}(\Gamma_i) \right\rangle = \int \mathcal{D}a \, e^{iS} \prod_i W_{R_i}(\Gamma_i) \tag{4.50}$$

will be diffeomorphism invariant: they will be pure numbers depending only on the way in which the various loops Γ_i are knotted and braided with respect to each other, and on the various representations R_i carried by each loop. Hence field theory can be used to obtain many classes of knot invariants [146] (but this will not be our interest here).

We will be interested in quantizing this theory on a space-time manifold, which we write as $M = \Sigma \times \mathbb{R}$, with $\mu = 0$ denoting the time direction. Selecting the gauge $a_0 = 0$, the canonical

Hamiltonian can be seen to be

$$\mathcal{H} = \frac{k}{4\pi} \operatorname{Tr} a_2 \partial_0 a_1 - a_1 \partial_0 a_2 - \mathcal{L} = 0. \tag{4.51}$$

The Hamiltonian vanishes, but the theory is not completely trivial since we must still identify its Hilbert space: *i.e.*, we must solve the gauge constraint. Switching to complex coordinates on Σ, we obtain a simple form if we work in the holomorphic gauge $a_{\bar{z}} = 0$. This field only appears in the action linearly, so the gauge constraint is that the t, z components of the field strength tensor vanish:

$$\partial_i a_j^\alpha - \partial_j a_i^\alpha + f^{\alpha\beta\gamma} a_i^\beta a_j^\gamma \tag{4.52}$$

where $i, j = t, z$. This may be solved by letting $a_i^\alpha = (\partial_i g) g^{-1}$ where g takes values in G. Imposing this constraint on the action gives

$$S = \frac{k}{4\pi} \int_\Sigma \operatorname{Tr} (\partial g^{-1})(\overline{\partial} g) + \frac{k}{12\pi} \int_M \epsilon^{\mu\nu\rho} \operatorname{Tr} (\partial_\mu g) g^{-1} (\partial_\nu g) g^{-1} (\partial_\rho g) g^{-1}. \tag{4.53}$$

Here the second term, despite appearances, only takes values on the boundary of M. This is the (chiral) Wess-Zumino-Witten (WZW) action [148], defined on the $2 + 0$ dimensional manifold Σ. The ground-state wavefunction of the CS theory is the partition function of the WZW model, and likewise if some Wilson lines were present which terminated on Σ we would obtain a conformal block of corresponding operators in the WZW model.

This result is important for many reasons, but perhaps chief among them is that all rational conformal field theories can be obtained as a WZW model for some suitably chosen group G. Frequently one needs to quotient out by a subgroup of G; in this way one obtains what are called coset models via the GKO construction [149]. For example, the Virasoro unitary models of section 4.2.3 may be obtained as the series of cosets $su(2)_k \oplus su(2)_1 / su(2)_{k+1}$, $k \geq 1$.

While this provides a method of constructing all known rational conformal field theories, it is important to note that many features of the theory are difficult to compute in this representation, such as singular vectors and correlation functions.

4.2.5 Superconformal symmetry

One possible symmetry we can add to the conformal group is supersymmetry [69, 150], which has the effect of exchanging bosons and fermions. This may be formulated in terms of theories living on supermanifolds, where some of the coordinates are anticommuting Grassman numbers θ. The stress tensor T_B will then have a conserved fermionic counterpart T_F, which we refer to as the supercurrent.

When we perform radial quantization, there are two possible boundary conditions for fermionic fields (such as T_F). This means the Hilbert space of this theory separates into two subspaces: the Neveu-Schwarz subspace, with periodic boundary conditions, and the Ramond subspace with antiperiodic boundary conditions. Note that our theory must include both types of boundary conditions, since it will turn out to contain operators (called spin fields) which switch between them.

Repeating the arguments of section 4.2.1, we may obtain the mode expansion of the stress tensor and supercurrent,

$$T_F(z) + \theta T_B(z) = \sum_{r,n} \frac{1}{2} \frac{G_r}{z^{r+3/2}} + \frac{\theta L_n}{z^{n+2}} \tag{4.54}$$

and obtain

$$\begin{aligned} [L_m, L_n] &= (m-n)L_{m+n} + \frac{\hat{c}}{8} m(m^2-1)\delta_{m+n,0}, \\ [L_m, G_r] &= \left(\frac{m}{2} - r\right) G_{m+r}, \\ \{G_r, G_s\} &= 2L_{r+s} + \frac{\hat{c}}{2}\left(r^2 - \frac{1}{4}\right)\delta_{r+s,0}. \end{aligned} \tag{4.55}$$

Here $\hat{c} = 3/2c$. Here we get the Neveu-Schwarz algebra for r, s half-integer and the Ramond algebra for r, s integer. The group of global superconformal transformations $OSP(2|1)$ is generated by $L_{\pm 1}, L_0, G_{\pm 1/2}$ and conjugates; we see that $G_{-1/2}^2 = L_{-1}$, meaning that a supersymmetry transformation is a "square root" of a translation in the plane. The analogous expression in the Ramond sector is $G_0^2 = L_0 - \frac{1}{16}\hat{c}$, the generator of translations on the cylinder.

As alluded to above, the operator-state correspondence must be modified due to boundary

condition issues. The vacuum $|0\rangle$, corresponding to the identity operator, belongs to the Neveu-Schwarz sector, and superfields correspond to states in this sector in the usual way: $|h\rangle = \Phi(0,0)|0\rangle$. In the Ramond sector, because G_0 commutes with L_0, any state is doubled: we may define $|h^-\rangle$ and $|h^+\rangle = G_0|h^-\rangle$. These states are obtained from the vacuum by the action of boundary-condition-changing spin fields: $|h^\pm\rangle = \theta^\pm(0)|0\rangle$. These pairs of states are degenerate (and hence supersymmetry is unbroken) if there exists a spin field in the theory with $h = \frac{1}{16}\hat{c}$.

Note that the fermionic parts of the superfields are nonlocal with respect to the spin fields, and the $+$ and $-$ components of the spin fields are nonlocal with respect to each other, due to square root branch cuts appearing in the corresponding OPEs. We can obtain a theory in which all fields are mutually local either by only considering the superfields, or by taking the part of the Neveu-Schwarz and Ramond sectors with the Witten operator $\Gamma = (-1)^F = 1$, where F is the fermion number operator. Such a theory is known as a spin model.

We may construct minimal models of the superconformal algebra just as we did for the Virasoro algebra; such models are denoted $SM(p,p')$ and are only unitary for $p' = p+2$.

4.3 Clustered states, residues and filtrations

In this section I present some existing work using new terminology and also give my new results. We will be concerned with generalizing Haldane's pseduopotential argument, which we accomplish by considering a more general projection (or "special") Hamiltonian whose terms consist of projectors that cause the wavefunction to vanish as r powers of separation whenever any $k+1$ particles approach the same spot.

4.3.1 k-body pseudopotential formalism

The appropriate few-body generalization of Haldane's argument was investigated by Simon, Rezayi and Cooper in [64]. Interactions involving more than two particles (which cannot be reduced to a set of two-particle interactions) may appear unphysical, but they may arise as effective couplings obtained by integrating out virtual processes, and they may also be added to the system explicitly, for example in the case of quantum Hall states realized as cold atomic gases in optical

traps. Alternatively, we may simply decide to view these Hamiltonians as existence proofs for particular states of matter.

We recall from the discussion in section 4.1.3 that Haldane's argument began by expanding the relative part of the two-particle wavefunction in an appropriate basis, which was simply the set of projection operators onto translationally invariant polynomials of degree (relative angular momentum) L. Simon *et. al.* proceed by doing the same for a $k+1$-body wavefunction.

Again, a basis state of relative angular momentum r will be a homogenous polynomial of degree r in $k+1$ variables. An important difference from Haldane's case is that k and r by themselves do not always uniquely determine the polynomial: there may be a finite-dimensional subspace for a given particle number and relative angular momentum. In order to determine the dimension and properties of this space, I first need to summarize a few facts from the theory of symmetric polynomials.

Symmetric polynomials

The starting point for the study of symmetric polynomials (that is, polynomials that are invariant under any permutation of their arguments) is the set of k elementary symmetric polynomials in k variables [151]:

$$\begin{aligned} e_1(z_1,\ldots,z_k) &= \sum_{j=1}^{k} z_j; \\ e_2(z_1,\ldots,z_k) &= \sum_{1\le j_1\le j_2\le k} z_{j_1} z_{j_2}; \\ e_3(z_1,\ldots,z_k) &= \sum_{1\le j_1\le j_2\le j_3\le k} z_{j_1} z_{j_2} z_{j_3}; \\ &\cdots \\ e_k(z_1,\ldots,z_k) &= \prod_{m=1}^{k} z_m; \end{aligned} \tag{4.56}$$

These arise, for example, in the expansion of the Laughlin quasihole factor (which serves as the generating function for the $\{e_i\}$)

$$\prod_{i=1}^{k}(z_i - w) = (-w)^k + (-w)^{k-1}e_1 + (-w)^{k-2}e_2 + \cdots + e_k, \tag{4.57}$$

where I have suppressed the arguments of the $\{e_i\}$. The fundamental theorem of symmetric polynomials [151] states that any symmetric polynomial may be expressed as a linear combination of products of the $\{e_i\}$.

In what follows, I define a partition $\lambda \vdash r$, $|\lambda| = k$ to be a nonincreasing sequence of integers $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k \geq 0$ such that $\sum_{i=1}^{k} \lambda_i = r$. I find it useful to define the occupation numbers $\mu_j(\lambda)$ as the number of parts of λ equal to j; then $\sum_j j\mu_j(\lambda) = r$.

With these definitions, the fundamental theorem of symmetric polynomials states that any homogenous symmetric function ψ_r with fixed number of variables k and homogeneous degree r may be written as

$$\psi_r(z_1, \ldots, z_k) = \sum_{\substack{\lambda \vdash r \\ \lambda_1 \leq k}} c_\lambda \prod_{j=1}^{k} e_j^{\mu_j(\lambda)}(z_1, \ldots, z_k) \equiv \sum_{\substack{\lambda \vdash r \\ \lambda_1 \leq k}} c_\lambda e_\lambda(z_1, \ldots, z_k), \tag{4.58}$$

for some constants $\{c_\lambda\}$ (in addition to a constant term, which I neglect in what follows). We must constrain the largest part of λ to be less than or equal to k since there are only k elementary symmetric polynomials in this many variables.

From this theorem, we see that the dimension of this space $D_{k,r}$ is equal to the number of integer partitions of r with no part larger than k. This function cannot be expressed in closed form, but I may easily give a generating function for it as

$$Z(q, w) \equiv \sum_{k,r=0}^{\infty} q^k w^r D_{k,r} = \sum_{i=0}^{\infty} w^i \prod_{j=1}^{i} \frac{1}{1 - q^j}. \tag{4.59}$$

Since we're essentially only doing linear algebra here, we should feel free to choose whatever basis is most convenient for the problem at hand. Another basis which will come in handy later is

the set of power sum symmetric functions,

$$p_j(z_1, \ldots, z_k) = \sum_{i=1}^{k} z_i^j \quad \text{for } 1 \leq j \leq k \tag{4.60}$$

from which we can form linear combinations by raising successive power sums to occupation numbers of some partition, as we did in (4.58).

Translationally-invariant symmetric polynomials

Having characterized the homogenous symmetric polynomials, we now turn our attention to the subspace of these which is translationally invariant and hence forms an appropriate basis for the relative part of the few-body wavefunctions. Any such function must be expressible solely in terms of coordinates relative to the center of mass:

$$\tilde{z}_i \equiv z_i - \frac{1}{k}\sum_{j=1}^{k} z_j, \quad \text{for } 1 \leq i \leq k. \tag{4.61}$$

(This is, of course, now an overcomplete set of coordinates). We can now make use of the fundamental theorem of symmetric polynomials as before, and determine that any function symmetric in the $\{\tilde{z}_i\}$ must take the form of sums of products of $e_k(\tilde{z}_1, \ldots, \tilde{z}_k)$ as before, and we may index those products by integer partitions. However, we have $e_1(\tilde{z}_1, \ldots, \tilde{z}_k) = 0$ for any k, so we now have the additional constraint that the partition cannot contain any parts equal to one. For $k > 1$, the other $e_k(\tilde{z}_1, \ldots, \tilde{z}_k)$ remain nonvanishing and linearly independent, which was proved in [64].

Therefore, we have the basis

$$\tilde{e}_\lambda(z_1, \ldots, z_k) = \sum_{j=2}^{k} e_j^{\mu_j(\lambda)}(\tilde{z}_1, \ldots, \tilde{z}_k), \tag{4.62}$$

for admissible partitions λ, and any translationally invariant symmetric function may be written as a combination

$$\psi_r(z_1, \ldots, z_k) = \sum_{\substack{\lambda \vdash r \\ \lambda_1 \leq k \\ \lambda_k > 1}} c_\lambda \tilde{e}_\lambda(z_1, \ldots, z_k). \tag{4.63}$$

The dimension of the translationally-invariant subspace is given by the number of partitions of L into parts no larger than k and greater than one. The generating function for the dimensions of the subspaces is then

$$\tilde{Z}(q,w) \equiv \sum_{k,r=0}^{\infty} q^k w^r \tilde{D}_{k,r} = \sum_{i=0}^{\infty} w^i \prod_{j=2}^{i} \frac{1}{1-q^j}. \tag{4.64}$$

Some values of $D_{k,r}$ and $\tilde{D}_{k,r}$ are tabulated in figure 4.1 for the reader's inspection. The Haldane pseudopotentials correspond to the $k=2$ row of the $\tilde{D}_{k,r}$ table: the entries are not all equal to one since we have restricted ourselves to bosonic statistics.

	$r=$	0	1	2	3	4	5	6	7	8	9	10	11	12
$D_{k,r}:$	$k=2$	1	1	2	2	3	3	4	4	5	5	6	6	7
	$k=3$	1	1	2	3	4	5	7	8	10	12	14	16	19
	$k=4$	1	1	2	3	5	6	9	11	15	18	23	27	34
	$k=5$	1	1	2	3	5	7	10	13	18	23	30	37	47
$\tilde{D}_{k,r}:$	$k=2$	1	0	1	0	1	0	1	0	1	0	1	0	1
	$k=3$	1	0	1	1	1	1	2	1	2	2	2	2	3
	$k=4$	1	0	1	1	2	1	3	2	4	3	5	4	7
	$k=5$	1	0	1	1	2	2	3	3	5	5	7	7	10

Figure 4.1: Tabulation of the dimensions $D_{k,r}$ of the space of symmetric polynomials in k variables that are homogenous of degree r, and of the subspace of those polynomials which are translationally invariant, which has dimension $\tilde{D}_{k,r}$.

The elements of this basis are orthogonal in the sense of polynomials (where we consider each distinct monomial $z_1^{n_1} z_2^{n_2} \cdots z_N^{n_N}$ in the polynomial's expansion to be a distinct basis vector), but we also have a geometric product with integration measure given by the geometry of the manifold the wavefunction is defined on. This inner product is given by

$$\langle \phi | \psi \rangle = \int d\mu_k \, \phi(z_1^*, \ldots, z_k^*) \psi(z_1, \ldots, z_k) \tag{4.65}$$

with the measure in the plane, for example, being given by

$$d\mu_k = \prod_{i=1}^{k} e^{-|z_i|^2/2} dz_i \, dz_i^*. \tag{4.66}$$

Since the polynomials do not know about this measure, in general they will not be orthogonal

with respect to it (although they will continue to be linearly independent as long as the measure is nonsingular.) I could orthogonalize them by the standard Gram-Schmidt algorithm, but since we wish to discuss different geometries, I will put off committing to a given inner product for as long as possible.

Geneneralized hard-core Hamiltonians

I have now defined the few-particle analogues of Haldane's projection operators. We can generalize the next step by considering hard-core cluster Hamiltonians, which for fixed k involve all projectors up to a certain value of r. In effect, these Hamiltonians will forbid $k+1$ particles to come together with relative angular momentum less than r — equivalently, the ground state wavefunction must vanish at least as fast as r powers of the separation when any $k+1$ particles are brought to the same point. From counting the overall degree the resulting wavefunction must have, it follows that the zero-energy states of these Hamiltonians describe quantum Hall states at filling fraction $\nu = k/r$ (for states of bosonic particles).

As noted above, the Laughlin states correspond to the $k = 2$ series. The Moore-Read state [119] is given by $k = 3, r = 2$, and the Read-Rezayi series [62] is given by $r = 2$. A number of other series for $k = 3$ and $r = 3$, $r = 4$ were conjectured by Simon, Rezayi and Cooper [51].

This progress is encouraging, but there are many familiar and seemingly simple CFTs which are missing from this series. We now arrive at the central question motivating this chapter: given a CFT whose blocks define a set of CFT states, can we obtain these blocks as polynomials subject to a projection Hamiltonian? This question is too ambitious to treat in full generality (as of this writing), so in what follows I will consider a modification of the hard-core formalism: for values of k, r which define a subspace with $\tilde{D}_{k,r} > 1$, I examine one particular Hamiltonian that projects out all states with $r' < r$ and *all but one* of the states at k, r. A different Hamiltonian of this same type was studied in [63]. Hamiltonians of this type will have a number of continuous parameters corresponding to the allowed state in this space. I now introduce a powerful and simple technique for studying such Hamiltonians and (perhaps more importantly) generating a basis of wavefunctions for their excited states.

4.3.2 The filtration method

The filtration method was introduced by Read [65, 66] in the context of counting excitations of the parafermion states [62], but it has general applicability to any clustered state arising from a special Hamiltonian.

As described above, we consider zero-energy states of a projection Hamiltonian requiring $k+1$ particles to vanish as r powers when they are brought together. I let $V_{N,M,k}$ denote the space of symmetric functions of N variables of degree at most M in each variable which vanish when $k+1$ particles are brought together.

We first consider the behavior of an element of $V_{N,M,k}$ when only k of its coordinates are identified. If it vanishes, this function also belongs to $V_{N,M,k-1}$. If not, we can define

$$C_1^k \Psi_{N,M,k}(z_1, \ldots, z_N) = \psi(Z_1)\Psi_{N-k,M,k-1}(z_{k+1}, \ldots, z_N). \tag{4.67}$$

where I use the notation $\Psi_{N,M,k}$ for members of the set $V_{N,M,k}$ and I define the clustering operator C_1^k which identifies the first k coordinates: $z_1 \mapsto Z_1, \ldots z_k \mapsto Z_1$. We refer to $\Psi_{N-k,M,k-1}$ as the "residue" of this procedure (following the terminology of [66], in which this procedure was defined on wavefunctions with the Laughlin factors omitted and ψ was a singular function).

We may now repeat the procedure on $\Psi_{N-k,M,k-1}(z_{k+1}, \ldots, z_N)$, since it remains a symmetric function, and so on, stopping only when we obtain a function that vanishes (or we run out of coordinates). This may be described in terms of clustering operators C_m^k which map the first m groups of k particles to identical values $Z_1, \ldots, Z_m$. We can do this without loss of generality since the $\{C_m^k\}$ act on symmetric functions, mapping them to linear combinations of other symmetric functions of the unpaired zs and the cluster coordinates $\{Z_i\}$. If we denote the space of symmetric polynomials in m variables as Λ_m, we have $C_m^k : \Lambda_N \to \Lambda_{N-mk} \otimes \Lambda_m$.

It may happen that the wavefunction in question vanishes when some of the particles are clustered, *i.e.* the wavefunction is a member of the kernel of some C_m^k. If this is the case, it certainly vanishes if we cluster more particles, so $\ker C_{m-1}^k \subseteq \ker C_m^k$. If we start off by acting on states that are a member of the ideal $I_N \subset \Lambda_N$ of states allowed by the Hamiltonian, we can define the

spaces $F_m^k = \ker C_m^k \cap I_N$. Since the kernels have a nested structure, this gives a filtration

$$F_0 = \{0\} \subseteq F_1 \subseteq F_2 \subseteq \cdots \subseteq F_{N/k} \subset I_N = F_{N/k+1}. \tag{4.68}$$

We now invoke details of the specific Hamiltonian and the behavior it allows. We first streamline notation by numbering our particles according to what cluster they will be in: we label them as $\{z_{i,j}\}$, with $1 \leq i \leq N/k$ and $1 \leq j \leq k$; then under action of the C^ks all coordinates with the same value of i will be mapped to Z_i. Then we may characterize the properties of polynomials in the kernel of C_m^k using the properties of our Hamiltonian: we must have

$$C_m^k : \Psi_{N,M,k} \to \prod_{i<i'} (Z_i - z_{i',j})^r \cdot \prod_{i<i';j} (Z_i - Z_{i'})^{rk} \cdot \prod_{i,l} (Z_i - w_l) \Psi_{N-mk,M-rm,k-1} \tag{4.69}$$

Here we have reused indices in factors separated by dots $\cdot$. The first factor must be present according to our Hamiltonian: if a $k+1$st particle is brought towards a cluster of k particles, any zero-energy state allowed by the Hamiltonian must vanish at least as fast as $(Z-z)^r$. The second factor also arises as a consequence of the projection requirement: we may bring two clusters by first making one at Z_1 and then assembling the other by successively bringing in k particles to a location Z_2 close to Z_1. Since the leading behavior as any one particle approaches Z_1 goes as $(Z_1 - z)^r$, we obtain k such factors when the second cluster is fully assembled. The third factor introduces auxiliary parameters w_l which will correspond to quasihole coordinates in a manner that will be made clear below.

At the end of this clustering, we are left with a residue that must necessarily vanish if any more coordinates are grouped into k-clusters (hence the third subscript $k-1$). The residue is a function of the $N - mk$ remaining unclustered coordinates, whose maximum degree is reduced by rm due to the coordinates appearing in the first factor. We may then inductively repeat this procedure on $\Psi_{N-mk,M-rm,k-1}$, and so forth, until the complete residue is built up.

The residue corresponds to the maximally clustered version of $\Psi_{N,M,k}$. The next step is to construct a basis of functions, each of which yields one of the set of residues we have just obtained.

4.3.3 Filtrations and symmetric polynomials

Particle-cluster behavior

The filtration is set up in terms of the clustering conditions: we require the wavefunction to vanish as r powers when $k+1$ particles are brought together, but the filtration recursively builds clusters which, additionally, do not vanish when the first $k, k-1, \ldots, 2$ particles are brought together. In what follows it will be useful to work in a basis of symmetric polynomials chosen to accommodate this behavior: we do this by first defining

$$\vartheta_\lambda(z_1, \ldots, z_k; z_{k+1}) = \prod_{j=1}^{k} (z_j - z_{k+1})^{\lambda_j} \tag{4.70}$$

for numbers $\lambda_1, \lambda_2, \ldots, \lambda_k$ such that $\sum_{i=1}^{k} \lambda_i = r$. Eq. (4.70) clearly is nonvanishing as $z_1, \ldots, z_k$ are brought together and vanishes as r powers when the $k+1$st particle is brought in. We obtain symmetric polynomials by explicitly symmetrizing with the operator $\mathcal{S}$ defined as follows:

$$\begin{aligned} \theta_\lambda(z_1, \ldots, z_{k+1}) &= \mathcal{S}\vartheta_\lambda(z_1, \ldots; z_{k+1}) \\ &= \frac{1}{(k+1)!} \sum_{\pi \in S_{k+1}} \vartheta_\lambda(z_{\pi(1)}, z_{\pi(2)}, \ldots; z_{\pi(k+1)}). \end{aligned} \tag{4.71}$$

Because of the explicit symmetrization, without loss of generality we may take λ to be a partition of r with $\leq k$ parts — this follows from only considering permutations which exchange the first k particles. Because ϑ, and hence θ, are translationally invariant by definition, the polynomials θ are members of the set of translationally-invariant symmetric polynomials defined in section 4.3.1. I claim that, for a given r, k, we can obtain a basis for that space as follows (see figure 4.2). Starting from an admissible partition λ which defines a basis polynomial in the sense of section 4.3.1, we take its transpose (interchanging rows and columns). For $k+1$ particles, the transpose partition will have $\leq k+1$ parts. If it has $k+1$ parts, we perform the additional step of appending the last row to the first row, otherwise we do nothing. This procedure defines a new partition $w(\lambda) = \lambda'$. I now show that the functions $\vartheta_{w(\lambda)}$ defined from the admissible partitions λ constitute a basis for the translationally-invariant polynomials.

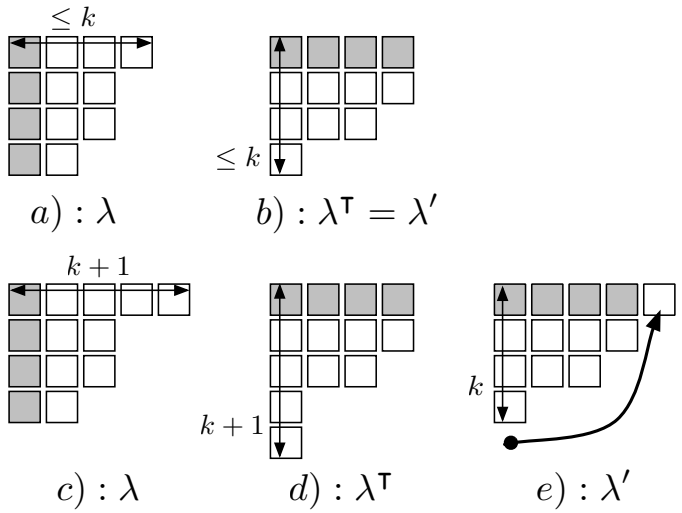

Figure 4.2: Example of various partitions involved in defining translationally-invariant symmetric polynomials, here in $k+1=5$ variables. Parts $a), b)$ illustrate the first case discussed in the text, in which the new partition $w(\lambda) = \lambda'$ is the transpose of the admissible partition λ. Parts $c), d)$ and $e)$ illustrate the extra rule we adopt when the width of λ is equal to the number of variables.

Proof of linear independence

It was proven in [64] that the polynomials $\tilde{e}_\lambda$ defined in (4.62) are a complete and linearly-independent basis when the partitions λ ranged over the set of partitions with width $\leq k+1$ and no parts of width one — "admissible" partitions in the terminology of section 4.3.1 and this section.

Elementary symmetric polynomials may be expanded in an alternative basis for symmetric polynomials (the monomial basis) as follows [151]:

$$e_\lambda = m_{\lambda^{\intercal}} + \sum_{\mu < \lambda^{\intercal}} M_{\lambda,\mu} m_\mu \tag{4.72}$$

where $<$ is a total ordering on the set of partitions which I will not need to specify (technically, it can be any ordering compatible in a certain sense with the dominance ordering) and $\lambda^{\intercal}$ is the transpose of λ. I have introduced the monomial symmetric polynomials

$$m_\lambda(z_1, \ldots, z_{k+1}) \propto \mathcal{S} z_1^{\lambda_1} z_2^{\lambda_2} \cdots z_{k+1}^{\lambda_{k+1}} \tag{4.73}$$

(the exact normalization is not important for my purposes). The fact that the change of basis matrix

$M_{\lambda,\mu}$ is upper triangular means that we have an automorphism $e_\lambda \mapsto m_{\lambda^\intercal}$ on the set of degree-r symmetric polynomials in $k+1$ variables. A similar map may be defined for the translationally-invariant subspace simply by treating the $\tilde{z}_i$ as the independent variables, rather than the z_i:

$$\tilde{e}_\lambda(z_1,\ldots,z_{k+1}) = e_\lambda(\tilde{z}_1,\ldots\tilde{z}_{k+1}) \mapsto m_{\lambda^\intercal}(\tilde{z}_1,\ldots\tilde{z}_{k+1}) \equiv \widetilde{\mathcal{S}}\tilde{z}_1^{\lambda_1^\intercal}\tilde{z}_2^{\lambda_2^\intercal}\cdots\tilde{z}_{k+1}^{\lambda_{k+1}^\intercal}. \tag{4.74}$$

Recall the definition (4.61) that $\tilde{z}_i = z_i - \frac{1}{k+1}\sum_{j=1}^{k+1} z_j$.

If the admissible partition λ has width $\leq k$, we may define another bijection to obtain our basis functions by shifting variables by $\tilde{z}_{k+1}$:

$$\begin{aligned} m_{\lambda^\intercal}(\tilde{z}_1,\ldots\tilde{z}_{k+1}) &\mapsto m_{\lambda^\intercal}(\tilde{z}_1-\tilde{z}_{k+1},\tilde{z}_2-\tilde{z}_{k+1},\ldots,0) \\ &= m_{\lambda^\intercal}(z_1-z_{k+1},z_2-z_{k+1},\ldots,0) = \theta_{\lambda'}(z_1,\ldots,z_{k+1}). \end{aligned} \tag{4.75}$$

where the function $\theta_{\lambda'}$ appearing in the second line is our goal, defined in (4.71). The fact that this is a bijection can be established by setting $\tilde{z}_{k+1}=0$ on both sides of the first line: $m_{\lambda^\intercal}(\tilde{z}_1,\ldots\tilde{z}_k) \neq m_{\mu^\intercal}(\tilde{z}_1,\ldots\tilde{z}_k)$ for $\lambda \neq \mu$.

If, on the other hand, λ has width $k+1$, the associated monomial is in the kernel of the mapping defined above, since we would encounter a factor of $(\tilde{z}_{k+1}-\tilde{z}_{k+1})^{\lambda_{k+1}^\intercal}$. We need to construct a partition λ' from $\lambda^\intercal$ which has $\leq k$ parts and is distinct from the transpose of any admissible partition. I choose to do this by appending the $k+1$th row of $\lambda^\intercal$ to its first row (figure 4.2, part e)): this results in a partition whose first two rows are unequal. The partitions $w(\lambda)=\lambda'$ obtained in this way are distinct from each other and those obtained in the previous paragraph: because admissible partitions λ have no part of width one, the first two rows of their transpose of their transpose have the same width. Again, the fact that these are distinct can be shown by setting $\tilde{z}_{k+1}$ to zero and comparing the monomial $m_{\lambda'}(\tilde{z}_1,\ldots\tilde{z}_k)$.

The preceding argument establishes that the mapping $w:\lambda\mapsto\lambda'$ is a bijection from the set of admissible partitions, and the mapping $\tilde{e}_\lambda \mapsto \theta_{w(\lambda)}$ is a bijection on the space of translationally-invariant polynomials. Note that I do not need to orthogonalize this basis, for reasons explained in [64]: the remainder of our discussion can be conducted without having to introduce an inner

product.

Cluster-cluster behavior

The polynomials θ_λ offer a complete description of any allowed behavior when one particle approaches a cluster of particles. Specifying this behavior also induces an interaction between clusters, since clusters are composed of particles, and this wavefunction can be built out of factors of ϑ_λ as follows. The most general case to consider concerns two clusters of $k_1 \geq k_2$ particles. I label coordinates by their cluster $i = 1, 2$ and their number within the cluster $j = 0, \ldots; k_i - 1$ as $\{z_{i,j}\}$. Then I define the cluster-cluster coupling

$$\begin{aligned} \overline{\theta}_\lambda(\{z_{1,i}\}, \{z_{2,j}\}) &= \mathcal{S}\overline{\vartheta}_\lambda(z_{1,0}, \ldots, z_{1,k_1-1}; z_{2,0}, \ldots, z_{2,k_2-1}) \\ &\equiv \mathcal{S} \prod_{m=0}^{k_2-1} \vartheta_\lambda(z_{1,m}, \ldots, z_{1,k_1-1+m}; z_{2,m}), \end{aligned} \tag{4.76}$$

where I take the j subscript of $z_{i,j}$ to be understood mod k_i; *i.e.* the interaction "wraps around" to earlier coordinates, as I illustrate in figure 4.3. This is necessary for the wavefunction to vanish in the same degree as every particle approaches a cluster. Also note that the product is cut off by the width of the narrower row. Finally, note that in what follows I use ϑ, θ to refer to particle-cluster couplings and $\overline{\vartheta}, \overline{\theta}$ to refer to the corresponding induced cluster-cluster coupling.

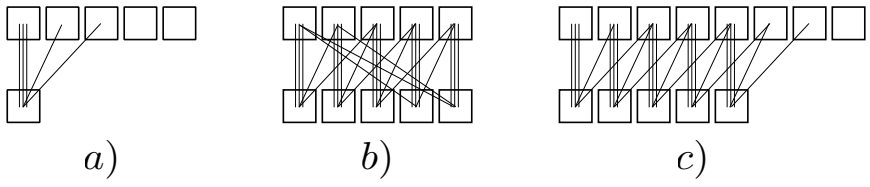

Figure 4.3: Schematic illustration of the pattern of connections involved in the various symmetric polynomials. Here the particle coordinates occupy boxes of a Young diagram and each diagonal or vertical line stands for a factor of $(z_{1,i} - z_{2,j})$. Figure $a)$ depicts $\vartheta_{\{3,1,1\}}$ for $k = 5$ particles. Figures $b)$ and $c)$ show $\overline{\vartheta}_{\{3,1,1\}}$ for the case of rows of equal and unequal width, respectively.

When we perform the clustering operation $z_{i,j} \to Z_i$, we have

$$C_1^{k_2} C_1^{k_1} \overline{\theta}_\lambda(Z_1, Z_2) \propto (Z_1 - Z_2)^{k_2 r}, \tag{4.77}$$

(recall that $\sum_i \lambda_i = r$), which establishes the minimal degree with which the wavefunction can vanish as two clusters are brought together.

In this chapter, I assume that all behavior is specified by the particle-cluster interactions (the functions ϑ_λ); as a consequence the cluster-cluster behavior (4.76) is built entirely out of these functions, and this means that when many clusters are present we only have two-body interactions between them, again induced by the particle-cluster interaction. The most general case is when the filtration procedure identifies clusters of $k_1 \geq k_2 \geq \cdots \geq k_\ell$ particles, by which I mean that in this maximally clustered state no more particles may be brought together without the wavefunction vanishing. I take these cluster numbers to define another integer partition $\kappa = \{k_1, k_2, \ldots, k_\ell\}$, with the total particle number $N = \sum_i k_i$. Then the couplings between these clusters are given by

$$\begin{aligned}\Psi^\kappa_\lambda(z_{1,0},\ldots,z_{1,k_1-1};z_{2,0},\ldots,z_{2,k_2-1};\ldots;z_{\ell,0},\ldots,z_{\ell,k_\ell-1}) \\ &\equiv \mathcal{S}\prod_{1\leq i<j\leq \ell}\overline{\vartheta}_\lambda(z_{i,0},\ldots,z_{i,k_i-1};z_{j,0},\ldots,z_{j,k_j-1}) \\ &= \mathcal{S}\prod_{1\leq i<j\leq \ell}\prod_{m=0}^{k_j-1}\vartheta_\lambda(z_{i,m},\ldots,z_{i,k_i-1+m};z_{j,m}).\end{aligned} \tag{4.78}$$

When the function Ψ^κ_λ is maximally clustered by sending $z_{i,j} \mapsto Z_i$, the behavior we obtain is

$$C^\kappa \Psi^\kappa_\lambda(Z_1,\ldots,Z_\ell) \propto \prod_{1\leq i<j\leq \ell}(Z_i - Z_j)^{k_j r} \tag{4.79}$$

As mentioned above, we may obtain a space V^κ_λ of symmetric polynomials with this behavior by multiplying Ψ^κ_λ by any symmetric function of the ℓ center of mass coordinates $\frac{1}{k_i}\sum_j z_{i,j}$ of the clusters. In the following sections I will use this functional basis to build sets of clustered wavefunctions describing ground and excited states of various quantum Hall systems.

In the following sections we will need to consider expressions that are more elaborate than Ψ^κ_λ, and displaying the explicit coordinate dependence of each function becomes cumbersome. In order to simplify our expressions, I add superscripts to $\overline{\vartheta}$ and ϑ which identify the row (cluster) that they operate on:

$$\overline{\vartheta}^{(i,j)}_\lambda \equiv \overline{\vartheta}_\lambda(z_{i,0},\ldots,z_{i,k_i-1};z_{j,0},\ldots,z_{j,k_j-1}). \tag{4.80}$$

$$\vartheta_\lambda^{(i)}(X) \equiv \prod_{z_{i',j}\in X} \overline{\vartheta}_\lambda(z_{i,j},\ldots,z_{i,k_i-1+j};z_{i',j}); \tag{4.81}$$

Here X stands for any set of coordinates in κ and, again, the column index in row i is understood mod k_i. In this notation, (4.78) may be abbreviated as

$$\Psi_\lambda^\kappa = \mathcal{S}\prod_{1\le i<j\le \ell} \overline{\vartheta}_\lambda^{(i,j)} = \mathcal{S}\prod_{1\le i<j\le \ell} \vartheta_\lambda^{(i)}(\{z_j\}). \tag{4.82}$$

where I introduce the additional notation that $\{z_i\}$ (with the column index suppressed) refers to the set of all coordinates in row i.

4.3.4 Superconformal theories at generic c

As mentioned above, a number of generalized ground states have been worked out or postulated for various values of k and r in [51, 62]. As a less trivial example, I am led to examine the case where there are multiple states at fixed angular momentum, and the Hamiltonian projects out only some of these (and all states with less angular momentum). A simple case to consider is that of three particles at degree six, which has two linearly-independent basis states.

This choice is motivated by the recent result of Simon [52], who found a Read-Rezayi-like form for the correlator of an arbitrary number of superconformal currents: defining

$$C_M = \langle G(z_1)G(z_2)\cdots G(z_M)\rangle\,, \tag{4.83}$$

the use of the superconformal OPEs (4.55) and the conformal Ward identity (4.27) implies

$$C_M = \sum_{i=2}^{M}(-1)^i\left[\frac{\frac{2}{3}c}{(z_1-z_i)^3} + \frac{2}{z_1-z_i}\sum_{j\neq 1,i}\left(\frac{3/2}{(z_i-z_j)^2} - \frac{1}{z_i-z_j}\frac{\partial}{\partial z}\right)\right]C_{M-2}(\hat{1},\hat{i}), \tag{4.84}$$

where $C_{M-2}(\hat{1},\hat{i})$ is the correlator with the insertions at z_1 and z_i removed. Simon's result was

that this recursion could be solved in closed form for arbitrary c: defining

$$\chi_4(z_1, z_2; z_3, z_4) = A z_{13}^3 z_{24}^3 z_{14}^3 z_{23}^3 + z_{13}^4 z_{24}^4 z_{14}^2 z_{23}^2 \tag{4.85}$$

where $z_{ij} = (z_i - z_j)$ and $A = c/3 - 1$ we have, up to a c-dependent normalization,

$$C_M \propto \prod_{(i,j)<(i',j')} (z_{i,j} - z_{i',j'})^{-3} \mathcal{S} \prod_{i>i'} \chi_4(z_{i,1}, z_{i,2}; z_{i',1}, z_{i',2}). \tag{4.86}$$

Up to the Jastrow factor, this is similar to the form of the functions θ_λ defined above. This suggests that we consider paired Hall states derived from superconformal field theories, where the electron operator is the product of G and a chiral boson. The ground state would then be given by (a suitable Jastrow factor times) (4.86). We might guess that such states would be given by a projection Hamiltonian which selects one of two ways for three particles to vanish as six powers when they approach each other, since this is the effect of the factors χ_4. Note that the allowed behavior of the polynomials is specified by c, the central charge of the SCFT. Since Simon has obtained the ground state, we now construct the excited states.

Filtration argument

Acting on the above hypothesis, we may write down a basis for all zero-energy excited states of such a Hamiltonian using a filtration argument [67]. Simon's wavefunction describes a set of clusters of two particles, such that when any additional particle approaches a pair the wavefunction vanishes as six powers of the separation. Here there are two admissible partitions, $\{3, 3\}$ and $\{2, 2, 2\}$. According to the procedure defined in the previous section, we have $w(\{2, 2, 2\}) = \{3, 3\}$ and $w(\{3, 3\}) = \{4, 2\}$ (where we used the extra rule in the latter case), so the two-dimensional space of potentially allowed behavior for three particles at degree six is spanned by

$$\theta_{\{3,3\}}(z_1, z_2, z_3) = \mathcal{S} z_{13}^3 z_{23}^3 = -5e_3^2(\tilde{z}_1, \tilde{z}_2, \tilde{z}_3) + 27e_2^3(\tilde{z}_1, \tilde{z}_2, \tilde{z}_3) \tag{4.87}$$

and

$$\theta_{\{4,2\}}(z_1,z_2,z_3)=\mathcal{S}z_{13}^4z_{23}^2=\frac{27}{2}e_3^2(\tilde{z}_1,\tilde{z}_2,\tilde{z}_3)-7e_2^3(\tilde{z}_1,\tilde{z}_2,\tilde{z}_3), \tag{4.88}$$

in my basis, where I write the polynomials in terms of coordinate differences $z_{ij}=(z_i-z_j)$ to make the translation invariance explicit. On the right-hand side I expressed these in the basis of (4.63). In my basis, Simon's cluster-cluster coupling (4.85) is

$$\overline{\chi}_A(z_1,z_2;z_3,z_4)=A\overline{\vartheta}_{\{3,3\}}+\overline{\vartheta}_{\{4,2\}} \tag{4.89}$$

with the particle-cluster coupling obtained from this relation by taking all particles in the second cluster except one to be very far away:

$$\begin{aligned}\chi_A(z_1,z_2;z_3)&=\lim_{z_4\to\infty}z_4^{-6}\overline{\chi}_A(z_1,z_2;z_3,z_4)\\&=A\vartheta_{\{3,3\}}(z_1,z_2;z_3)+\vartheta_{\{4,2\}}(z_1,z_2;z_3).\end{aligned} \tag{4.90}$$

When the first two particles in χ_A are paired we have

$$C_1^2\chi_A(Z_1,z_{2,0})=(1+A)(Z_1-z_{2,0})^6, \tag{4.91}$$

which vanishes as $A=-1$ which corresponds to $c=0$ (recall $A=c/3-1$). When this happens, the three-particle interaction is reducible to the product of two-particle interactions only, and χ_A reduces to a Jastrow factor:

$$\mathcal{S}\chi_{-1}=z_{12}^2z_{13}^2z_{23}^2. \tag{4.92}$$

This means that at $A=-1$ we obtain the Laughlin state for bosons at $\nu=1/2$. This was discussed in section 4.1.3, so in what follows I assume $A\neq-1$.

When all the particles are maximally clustered, the considerations of section 4.3.2 dictate that we should have a residue of the form (4.69) with $r=6$ and $k=2$, *i.e.* we should have factors of the form

$$\prod_{i\leq N-2m}\prod_{k\leq m}(z_i-Z_k)^6\cdot\prod_{1\leq k<l\leq m}(Z_k-Z_l)^{12}, \tag{4.93}$$

for some maximal number m of clustered pairs of particles. In this and what follows I use dots $\cdot$ to separate terms in which indices are reused. We know that for $A \neq -1$ the polynomials in the kernel of C_1^2 should be D_2 times some polynomial vanishing when three particles coincide. However, the set of polynomials vanishing when three particles coincide is precisely the set of zero-energy states of the Moore-Read projection Hamiltonian [129], which defines a polynomial ideal I^{MR} (if a polynomial is in I^{MR}, so is the product of that polynomial and any symmetric polynomial). The same argument applies for any number of preexisting paired particles, so we see that the mth residue must actually be of the form

$$\prod_{2m<i<j} (z_i - z_j)^2 \cdot \prod_{2m<i} \prod_{k\leq m} (z_i - Z_k)^6 \cdot \prod_{k<l\leq m} (Z_k - Z_l)^{12} \cdot I^{\mathrm{MR}}_{N-2m} \otimes \Lambda_m. \tag{4.94}$$

Again, for this residue the maximally clustered state involves m pairs of particles with coordinates Z_i, leaving $N - 2m$ particles unpaired. We can also have an arbitrary charge excitation of the clusters, which is an element of the space of symmetric polynomials in the cluster coordinates (the Λ_m factor).

Basis states in the plane

Based on this residue, I now construct a basis for excitations following the arguments of the previous section, with each element of the basis corresponding to a unique residue. The structure of polynomials in I^{MR} was determined in [66, 129] (see also [65, 152, 153] for the generalization of the construction to the Read-Rezayi series): in my notation they involve a partition κ with n_1 rows of width two (the paired clusters) and n_2 rows of width one (the unpaired particles). In my terminology, particle coordinates are associated with entries of κ and labeled as $z_{i,j}$, with i labeling the cluster.

The projection Hamiltonian for the Moore-Read state [129] requires the state to vanish as at least two powers of separation when any three particles coincide. A maximally clustered state will then consist of $n_1 \leq N/2$ clusters of two particles at positions Z_i and $n_2 = N - 2n_1$ unpaired particles at positions $z_{i,0}$, such that the wavefunction vanishes if any other coordinates are forced to coincide. The residue obtained is of the form (4.69) with degree $r = 2$ and clusters of $k = 2$

particles:

$$C_{n_1}^2 \Psi^{\mathrm{MR}}_{\kappa;\lambda^{(1)},\lambda^{(2)}} = \prod_{1\le i<j\le n_1} (Z_i - Z_j)^4 \cdot \prod_{1\le i\le n_1<j\le n_1+n_2} (Z_i - z_{j,0})^2 \times e_{\lambda^{(1)}}(Z_1,\dots,Z_{n_1}) \cdot e_{\lambda^{(2)}}(z_{n_1+1,0},\dots,z_{n_1+n_2,0}). \quad (4.95)$$

Again, the essential part of the filtration argument is that the basis states are in one-to-one correspondence with the residues, and the set of residues completely enumerates the way a set of particles can behave subject to the constraints imposed by the Hamiltonian. The states are then given by

$$\Psi^{\mathrm{MR}}_{\kappa;\lambda^{(1)},\lambda^{(2)}} = \mathcal{S} \prod_{i<j\le n_1} \overline{\vartheta}^{(i,j)}_{\{1,1\}}(z_{i,0}, z_{i,1}; z_{j,0}, z_{j,1}) \cdot \prod_{i\le n_1<j} \vartheta^{(i)}_{\{1,1\}}(z_{i,0}, z_{i,1}; z_{j,0}) \times e_{\lambda^{(1)}}\left(\tfrac{1}{2}(z_{1,0}+z_{1,1}),\dots,\tfrac{1}{2}(z_{n_1,0}+z_{n_1,1})\right) \times e_{\lambda^{(2)}}(z_{n_1+1,0},\dots,z_{n_1+n_2,0}). \quad (4.96)$$

Here all rows are coupled to each other with $\overline{\vartheta}_{\{1,1\}}$, the only possible choice of translationally-invarant polynomial at degree $r = 2$, and the partitions $\lambda^{(i)}$ with $\lambda^{(i)}_1 \le n_i$ describe charge excitations in the clusters $(i = 1)$ and unpaired particles $(i = 2)$. When this state is maximally clustered, we recover (4.95).

We may now combine this description of I^{MR} with the form (4.94) of the residues for the superconformal state to obtain a set of basis polynomials in the plane. This is given by a partition $\kappa \vdash N$ with two blocks of width $r_1, r_2 = 2$ and length m_1, m_2 and one block of width $r_3 = 1$ and length m_3. I introduce an additional, superscripted index to the particle coordinates to label which

block they belong to — see part $a)$ of figure 4.4. The functions for this partition are [67]

$$\Psi^{\text{SCFT}}_{\kappa;\lambda^{(1)},\lambda^{(2)},\lambda^{(3)}}[A] = \mathcal{S} \prod_{i<j\leq m_1} \overline{\chi}^{(i,j)}_A \cdot \prod_{i\leq m_1} \chi^{(i)}_A(\{z^{(2)}, z^{(3)}\}) \cdot D_2(\{z^{(2)}, z^{(3)}\}) \times \prod_{m_1<i<j\leq m_2} \overline{\vartheta}^{(i,j)}_{\{1,1\}} \cdot \prod_{m_1<i\leq m_2} \vartheta^{(i)}_{\{1,1\}}(\{z^{(3)}\}) \cdot \prod_{m_2<i<j\leq m_3} \overline{\vartheta}^{(i,j)}_{\{1,1\}} \times e_{\lambda^{(1)}}(\{Z^{(1)}\}) \cdot e_{\lambda^{(2)}}(\{Z^{(2)}\}) \cdot e_{\lambda^{(3)}}(\{Z^{(3)}\}). \quad (4.97)$$

In order to keep this expression manageable, I use the abbreviations defined in (4.80), (4.81). The notation $\{z^{(\alpha)}\}$ refers to the set of all coordinates in block α, and I have introduced the cluster center of mass coordinates

$$Z^{(\alpha)}_i = \frac{1}{r_\alpha} \sum_{j=0}^{r_\alpha} z^\alpha_{i,j}. \quad (4.98)$$

The function D_2 is a discriminant or Jastrow factor, arising from the first factor in the residue (4.94):

$$D_n(X) = \prod_{\substack{z_i,z_j\in X \\ i<j}} (z_i - z_j)^n. \quad (4.99)$$

Physically, the particles in block one correspond to fully paired clusters, those in block two correspond to excited, "half-broken" pairs, and those in the third block to further excited, fully broken pairs. A schematic depiction of the factors of (4.97) is given in part $b)$ of figure 4.4. Note that if $m_2 = m_3 = 0$ we recover Simon's result for the ground state, (4.86) (up to a Jastrow factor): in this case only the first product (over the $\overline{\chi}_A$s) contributes.

Character formulae

We can now easily compute the character function counting number of states of the form (4.97) for each value of total angular momentum. As I have noted above, this is a sum over all basis polynomials of q raised to the overall degree of the polynomial. The degree of the factor in the first

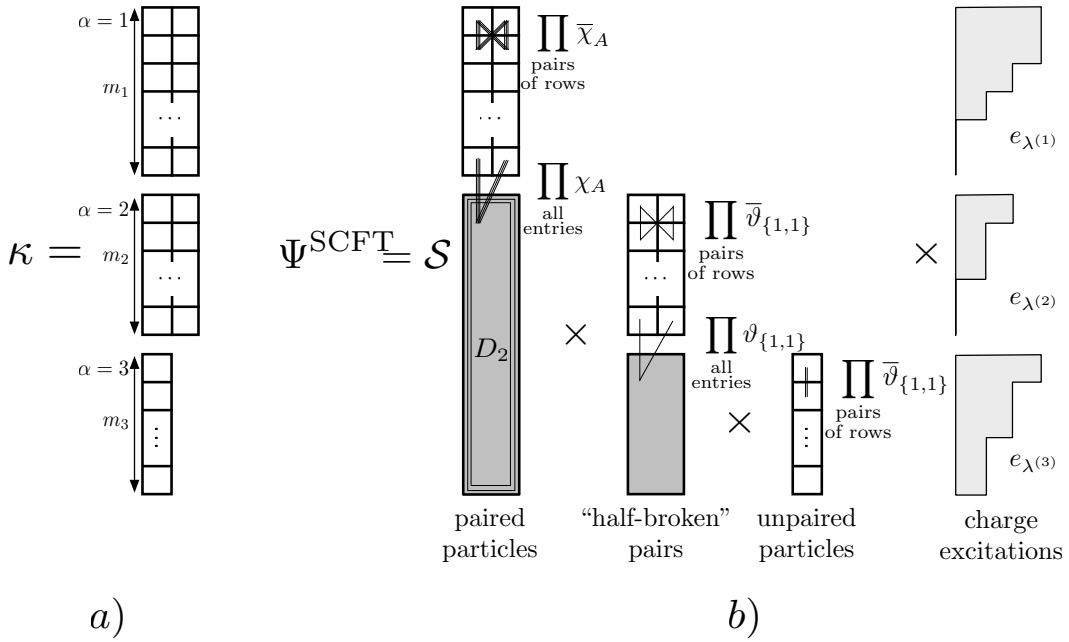

Figure 4.4: Diagram of the organization of factors in the state (4.97). Part $a)$ labels the organization of the coordinates, each of which occupies a box in the diagram of κ. Part $b)$ depicts each product appearing in (4.97): here for clarity I only draw the first factor in each product, using thin lines to stand for factors of coordinate differences $(z_{i,j}^{(\alpha)} - z_{i',j'}^{(\alpha')})$.

two lines of (4.97) is

$$d^{\mathrm{SCFT}}(\kappa) = 6m_1(m_1-1) + 6m_1(2m_2+m_3) + (2m_2+m_3)(2m_2+m_3-1) \\ + 2m_2(m_2-1) + 2m_2m_3 + m_3(m_3-1) \quad (4.100)$$

which means the overall character is

$$\mathrm{ch}\, I_N^{\mathrm{SCFT}} = \sum_{\substack{m_1,m_2,m_3: \\ 2m_1+2m_2+m_3=N}} \frac{q^{d^{\mathrm{SCFT}}(\kappa)}}{(q)_{m_1}(q)_{m_2}(q)_{m_3}}. \quad (4.101)$$

Here the factors of $1/(q)_{m_i} = \prod_{j=1}^{m_i} 1/(1-q^j)$ arise from counting the charge excitations, *i.e.* those partitions $\lambda^{(i)}$ with $\lambda_1^{(i)} \leq m_i$ (recall eq. (4.59)).

Connection with conformal field theory

To make contact with conformal field theory, I take the limit $N \to \infty$ of (4.101). Because

$$d^{\text{SCFT}} = \frac{3}{2}N(N-2) + 2m_2 + m_3 + \frac{1}{2}m_3^2 \tag{4.102}$$

the ground state will be obtained by taking $m_2 = 0$ and $m_3 = 0$ or 1 depending on the parity of N: we put as many particles as possible into the first block in order to minimize the overall angular momentum (degree). This is referred to as "filling the Bose sea" in [152]. Then

$$\lim_{N\to\infty} q^{-\frac{3}{2}N(N-2)} \text{ch } I_N^{\text{SCFT}} = \frac{1}{(q)_\infty} \sum_{\substack{m_2,m_3 \\ (-1)^{m_3}=(-1)^N}} \frac{q^{2m_2+m_3+\frac{1}{2}m_3^2}}{(q)_{m_2}(q)_{m_3}}, \tag{4.103}$$

where the sum is restricted to either even or odd values of m_3 depending on the parity of N: this, in effect, selects the vacuum sector of the theory.

This is the character of the Kac module for a superconformal field theory at a generic value of c [67] (times a factor of $1/(q)_\infty$ arising from the charge excitations). In other words, we have constructed the Verma module obtained by removing the vacuum sector singular vectors at levels 1/2 and 1, corresponding to the invariance of the vacuum under the global superconformal transformations (section 4.2.5). In particular, when A corresponds to the central charge of a superconformal minimal model our module is still reducible: our construction (and, by extension, our Hamiltonian) does *not* take into account these c-dependent, non-generic singular vectors. This should be surprising, as our projection Hamiltonian (choice of A in χ_A) *did* depend on c; on the other hand, eliminating the non-generic singular vectors would require that the spectrum of the Hamiltonian would depend in a very complicated, highly discontinuous way on the parameter c, which is unfeasible.

I note in passing that basis states and characters for the states on the sphere were also obtained in [67] (one makes the basis states homogenous in each z by introducing quasihole factors); I do not discuss them here because at the time of writing I have not been able to advance beyond these results.

4.3.5 The $M(3,p)$ series

We have seen, as a consequence of our filtration argument, that the projection Hamiltonian which eliminates three-particle states with angular momentum less than six and allows one such state at angular momentum six only reproduces the vacuum module of a generic superconformal theory; *i.e.* one with no non-trivial singular vectors. To obtain one of the superconformal minimal models, we need to remove more states by adding more projection operators to the Hamiltonian. Unfortunately, determining these operators from the SCFT itself appears difficult, and as of this writing I have made limited progress in this area.

One point of comparison is provided by Fegin, Jimbo and Miwa [68], who studied the $M(3,p)$ series of Virasoro minimal models using arguments similar to the filtration method described above. The $M(3,p)$ series includes $M(3,8)$, which is isomorphic to the superconformal minimal model $SM(2,8)$ and so should be describable with the above technique. Because we are now working at fixed values of c, the functions ϑ coupling the blocks together will be derived from the four-point functions of the corresponding CFTs, which may be obtained as the solutions to hypergeometric differential equations [135]. In particular, for the quantum Hall state derived from the $M(3,p)$ model (by the addition of a chiral boson) we have [68]

$$\begin{aligned}\varphi_4[p] &\equiv \langle 0|\phi_{2,1}(z_1)\phi_{2,1}(z_2)\phi_{2,1}(z_3)\phi_{2,1}(z_4)|0\rangle \prod_{1\le i<j\le 4}(z_i-z_j)^{(p-2)/2} \\ &= (z_{14}z_{23})^{(p-2)}{}_2F_1\left[1-\frac{p}{3},2-p,2-\frac{2p}{3};\frac{z_{12}z_{34}}{z_{14}z_{23}}\right].\end{aligned} \tag{4.104}$$

Here $\phi_{2,1}$ is an operator in the $M(3,p)$ theory and the Jastrow factor $\prod(z_i-z_j)^{(p-2)/2}$ arises from the correlator of suitably chosen vertex operators of the chiral boson (and ensures that the overall correlator is nonsingular for any choice of coordinates). The Gauss hypergeometric function is

denoted by ${}_2F_1$. The function (4.104) has the first few special cases

$$\begin{aligned}
\varphi_4[4] &= 3\overline{\theta}_{\{1,1\}};\\
\varphi_4[5] &= 3\overline{\theta}_{\{2,1\}};\\
\varphi_4[7] &= 3\overline{\theta}_{\{3,2\}};\\
\varphi_4[8] &= \mathcal{S} - \tfrac{12}{7}\left(\overline{\vartheta}_{\{4,2\}} - \tfrac{11}{4}\overline{\vartheta}_{\{3,3\}}\right) \propto \mathcal{S}\,\overline{\chi}_{A=-11/4}.
\end{aligned} \tag{4.105}$$

The overall normalizations are not important; note also that the last line agrees with Simon's result (4.89) evaluated at $c = c_{3,8} = -21/4$ (recall that $A = c/3 - 1$).

The residues for the $M(3,p)$ series obtained in [68] are

$$\prod_{2m<i<j} (z_i - z_j)^2 \cdot \prod_{2m<i} \prod_{k\le m} (z_i - Z_k)^{p-2} \cdot \prod_{k<l\le m} (Z_k - Z_l)^{2(p-2)} \cdot I^{[3,p-3]}_{N-2m} \otimes \Lambda_m. \tag{4.106}$$

For $p = 8$ and $M(3,8) = SM(2,8)$, the first three factors are in agreement with the residue of the generic superconformal theory (4.94). The important difference arises in the ideal $I^{[3,5]}$ describing the behavior of the unpaired particles: this is *not* the ideal $I^{\mathrm{MR}} = I^{[3,4]}$ encountered above, but instead describes zero-energy states of the Gaffnian state [154], a nonunitary state constructed from $M(3,5)$; it can be derived from a Hamiltonian requiring $k+1 = 3$ particles to vanish as $r = 3$ powers as they come together. Again, this means that the residues are of the form (4.69) with degree $r = 3$ and clusters of $k = 2$ particles:

$$\begin{aligned}
C^2_{n_1} \Psi^{\mathrm{G}}_{\kappa;\lambda^{(1)},\lambda^{(2)}} = \prod_{1\le i<j\le n_1} (Z_i - Z_j)^6 \cdot \prod_{1\le i\le n_1<j\le n_1+n_2} (Z_i - z_{j,0})^3 \\
\times e_{\lambda^{(1)}}(Z_1,\ldots,Z_{n_1}) \cdot e_{\lambda^{(2)}}(z_{n_1+1,0},\ldots,z_{n_1+n_2,0}),
\end{aligned} \tag{4.107}$$

which implies a corresponding set of basis states

$$\Psi^{\mathrm{G}}_{\kappa;\lambda^{(1)},\lambda^{(2)}} = \mathcal{S} \prod_{i<j\leq n_1} \overline{\vartheta}^{(i,j)}_{\{2,1\}}(z_{i,0}, z_{i,1}; z_{j,0}, z_{j,1}) \cdot \prod_{i\leq n_1<j} \vartheta^{(i)}_{\{2,1\}}(z_{i,0}, z_{i,1}; z_{j,0}) \times e_{\lambda^{(1)}}\left((z_{1,0}+z_{1,1})/2, \ldots, (z_{n_1,0}+z_{n_1,1})/2\right) \times e_{\lambda^{(2)}}(z_{n_1+1,0}, \ldots, z_{n_1+n_2,0}). \quad (4.108)$$

via arguments given previously (see also [154], footnote 19). This result should be compared with the Moore-Read states and residues obtained above: since we now work at degree three instead of two, clusters must be coupled with the function $\overline{\vartheta}_{\{2,1\}}$, which is the unique function in my basis for $k=2, r=3$. Here $\{2,1\} = w(\{3\})$.

This means that in order to obtain the basis states of the $M(3,8)$ theory, we simply replace the Moore-Read states with the Gaffnian in (4.97):

$$\Psi^{\mathrm{M(3,8)}}_{\kappa;\lambda^{(1)},\lambda^{(2)},\lambda^{(3)}} = \mathcal{S} \prod_{i<j\leq m_1} \overline{\chi}^{(i,j)}_{-11/4} \cdot \prod_{i\leq m_1} \chi^{(i)}_{-11/4}(\{z^{(2)}, z^{(3)}\}) \cdot D_2(\{z^{(2)}, z^{(3)}\}) \times \prod_{m_1<i<j\leq m_2} \overline{\vartheta}^{(i,j)}_{\{2,1\}} \cdot \prod_{m_1<i\leq m_2} \vartheta^{(i)}_{\{2,1\}}(\{z^{(3)}\}) \cdot \prod_{m_2<i<j\leq m_3} \overline{\vartheta}^{(i,j)}_{\{1,1\}} \times e_{\lambda^{(1)}}(\{Z^{(1)}\}) \cdot e_{\lambda^{(2)}}(\{Z^{(2)}\}) \cdot e_{\lambda^{(3)}}(\{Z^{(3)}\}). \quad (4.109)$$

Again, the partition κ is the same as in the generic superconformal case (4.97) (two blocks of width two and height m_1, m_2 and one block of width one and height m_3); the only difference between this state and (4.97) occurs in the second line, where clusters in block two are coupled to each other using $\overline{\vartheta}_{\{2,1\}}$ instead of $\overline{\vartheta}_{\{1,1\}}$ and coupled to the unpaired particles using $\vartheta_{\{2,1\}}$ instead of $\vartheta_{\{1,1\}}$. The situation is depicted in figure 4.5.

I emphasize that the validity of the basis (4.109) is ensured by the fact that the functions are in one-to-one correspondence with the known residues (4.106) — beyond that, the specific details of which functions are used amount to an arbitrary choice which is up to us. Because the residues (4.106) are so similar to those from the generic superconformal case (4.94), the proof given in [67] carries through essentially unchanged: because of the pattern of couplings, the only nonvanishing

way to pair two particles is if they are in the same row of block $\alpha = 1$. Once these particles have been paired, we may strip off the remaining Jastrow factor D_2 and obtain a Gaffnian wavefunction in $I^{[3,5]}$, as described above. As always, charge excitations are in one-to-one correspondence with symmetric polynomials in the cluster center of mass coordinates.

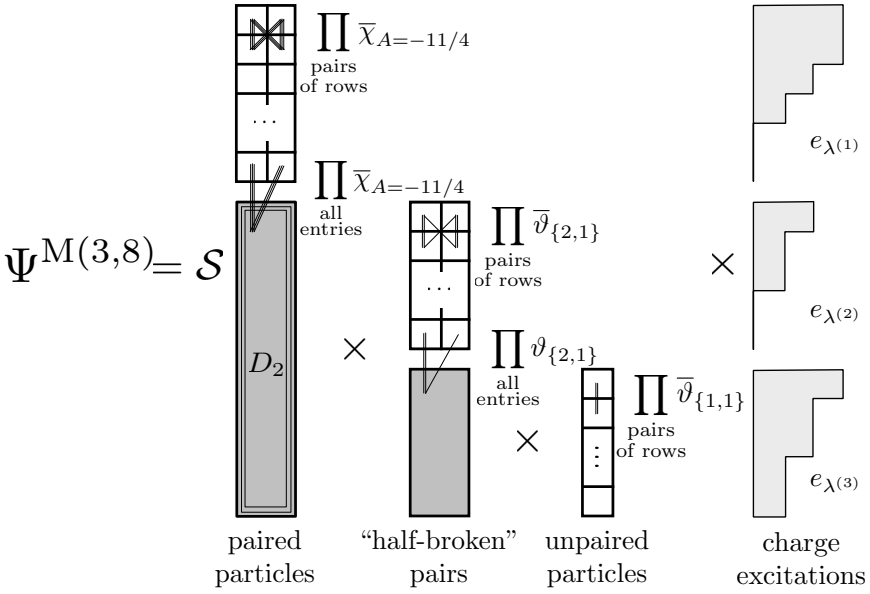

Figure 4.5: Diagram of the organization of factors in the state (4.109), which should be compared with figure 4.4 showing the analogous situation for the generic superconformal theory. Here for clarity I only draw the first factor in each product, using thin lines to stand for factors of coordinate differences $(z_{i,j}^{(\alpha)} - z_{i',j'}^{(\alpha')})$.

Characters

The replacement $\overline{\vartheta}_{\{1,1\}}, \vartheta_{\{1,1\}} \to \overline{\vartheta}_{\{2,1\}}, \vartheta_{\{2,1\}}$ changes the overall degree of the polynomials and hence the character; we now have

$$d^{\mathrm{M}(3,8)}(\kappa) = d^{\mathrm{SCFT}}(\kappa) + m_2 m_3 + m_2(m_2 - 1) \tag{4.110}$$

and the character at finite N is now

$$\mathrm{ch}\, I_N^{\mathrm{M}(3,8)} = \sum_{\substack{m_1,m_2,m_3: \\ 2m_1+2m_2+m_3=N}} \frac{q^{d^{\mathrm{M}(3,8)}(\kappa)}}{(q)_{m_1}(q)_{m_2}(q)_{m_3}}. \tag{4.111}$$

The $N \to \infty$ limit of this is

$$q^{-\frac{3}{2}N(N-2)}\text{ch } I_N^{\text{M(3,8)}} = \frac{1}{(q)_\infty} \sum_{\substack{m_2,m_3 \\ (-1)^{m_3}=(-1)^N}} \frac{q^{\frac{1}{2}m_3^2+m_2m_3+m_2^2+m_2+m_3}}{(q)_{m_2}(q)_{m_3}}, \tag{4.112}$$

which agrees with the fermionic characters for $M(3,8)$ (that is, $\chi_{1,1}^{(3,8)}(q)$ for N even and $\chi_{2,1}^{(3,8)}(q)$ for N odd) in the existing literature [155], so my basis contains the proper number of states and I conclude that I have accounted for all the singular vectors appearing in the vacuum sector of $SM(2,8)$.

Additional projection terms

What is the additional term we need to add to our Hamiltonian to obtain this basis? The constraint we need to impose was found at the operator level by Fiegin, Jimbo and Miwa [68] (their proposition 2.5), but really there is only one possible choice: namely, the allowed behavior must be of the form $\tilde{e}_2\varphi_4[p=8]$, otherwise we will no longer have a polynomial ideal. Specifically, the Hamiltonian we have permits one three-particle state at degree six. There is one possible charge excitation from that state to the space of states at degree eight (namely, multiplication by $\tilde{e}_2$), and this must be allowed otherwise we will not recover the proper contribution to the character for charge excitations (the factor of $1/(q)_\infty$ appearing in (4.103)). This is confirmed numerically: the allowed polynomial is

$$\mathcal{S} - \tfrac{5}{81}\vartheta_{\{5,3\}} - \tfrac{2}{81}\vartheta_{\{4,4\}} = 9\tilde{e}_3^2\tilde{e}_2 - \tilde{e}_2^4. \tag{4.113}$$

As I explained above, the state orthogonal to that depends on the geometry of the surface the wavefunctions are defined on; for the plane the state we must remove is

$$\mathcal{S}\,\tfrac{44}{9}\vartheta_{\{5,3\}} - \tfrac{43}{9}\vartheta_{\{4,4\}} = 11\tilde{e}_2^4 + 54\tilde{e}_3^2\tilde{e}_2. \tag{4.114}$$

I mention that a possible way to impose constraints of this form while maintaining the polynomial ideal structure would be by placing restrictions on how the clusters may approach each other. As we saw in section 4.3.3, any particle-cluster interaction induces a cluster-cluster inter-

action of Jastrow form, but clusters may potentially vanish faster than this; in addition, we might imagine adding multiple-cluster interactions. For example, we can also produce the $M(3,8)$ states by multiplying Ψ^{SCFT} by

$$\prod_{i,j}(Z_i^{(2)} - Z_j^{(2)})^2 \cdot \prod_{i,j}(Z_i^{(2)} - Z_j^{(3)}) \tag{4.115}$$

under the symmetrization, where, again, $Z_i^{(\alpha)}$ is the center of mass coordinate of the ith cluster in block α. This reproduces the character (4.111).

4.3.6 Remarks on the tricritical Ising model

In the previous section, we obtained states in the $SM(2,8)$ minimal model by adding extra couplings to the states of the generic superconformal theory (the replacement $\overline{\vartheta}_{\{1,1\}}, \vartheta_{\{1,1\}} \rightarrow \overline{\vartheta}_{\{2,1\}}, \vartheta_{\{2,1\}}$ implemented in (4.109)). One might guess that this would be a recipe for obtaining the states of any rational superconformal theory: in the generic states (4.97), the clusters and particles vanish in the lowest degree consistent with superconformal symmetry, and a specific superconformal minimal model would be obtained by adding extra couplings to make the wavefunction vanish faster in certain situations.

Fermionic characters

A promising model to test this hypothesis on is the tricritical Ising model, which is described by $M(4,5)$ or $SM(3,5)$ and is unitary. Importantly, unitarity means that, unlike $SM(2,8)$, this theory has a chance of describing a physical QHE state [53]. Again, we can gain some idea of which states need to be removed by inspecting fermionic characters. One form of the character for the vacuum sector of this model is [156]

$$\hat{\chi}_{1,1}^{(3,5)} = \sum_{n_1,n_2\geq 0} \frac{q^{\frac{1}{2}n_1^2+2n_2^2-n_1n_2}}{(q)_{2n_2}} \begin{bmatrix} n_2 \\ n_1 \end{bmatrix}_q , \tag{4.116}$$

which appears problematic, due to the presence of q-binomial terms $\begin{bmatrix} n_2 \\ n_1 \end{bmatrix}_q = \frac{(q)_{n_2}}{(q)_{n_1}(q)_{n_2-n_1}}$. My understanding at the time of writing is that factors of $1/(q)_n$ arise from charge excitations of a certain species of clusters, and I cannot currently obtain character functions with factors of $(q)_n$ in the numerator from counting basis states in the plane.

An alternate, equivalent form of the fermionic character may be derived by using this model's construction as the $(E_7^{(1)})_1 \otimes (E_7^{(1)})_1/(E_7^{(1)})_2$ coset model [157]:

$$\hat{\chi}_{1,1}^{(3,5)} = \sum_{\vec{m}\geq 0} \frac{q^{\vec{m}A\vec{m}}}{\prod_{i=1}^{7}(q)_{m_i}}, \tag{4.117}$$

where we sum over seven indices $\vec{m} = \{m_1, \ldots, m_7\}$ and the matrix A is the inverse Cartan matrix of E_7,

$$A = \begin{pmatrix} \frac{3}{2} & 1 & \frac{3}{2} & 2 & 2 & \frac{5}{2} & 3 \\ 1 & 2 & 2 & 2 & 3 & 3 & 4 \\ \frac{3}{2} & 2 & \frac{7}{2} & 3 & 4 & \frac{9}{2} & 6 \\ 2 & 2 & 3 & 4 & 4 & 5 & 6 \\ 2 & 3 & 4 & 4 & 6 & 6 & 8 \\ \frac{5}{2} & 3 & \frac{9}{2} & 5 & 6 & \frac{15}{2} & 9 \\ 3 & 4 & 6 & 6 & 8 & 9 & 12 \end{pmatrix}. \tag{4.118}$$

Here the vacuum sector (corresponding to the parity of N in the polynomial formalism) is selected by whether $m_1 + m_3 + m_6$ is even or odd, as can be seen by the presence of half-integer entries. This implies that the partition κ should now consist of *eight* blocks (including one for the "Bose sea" of unbroken clusters), three of which have odd width. The logical guess would be to have five blocks of width two and three of width one, which would be obtained from subdividing the three blocks of the partition κ used in the generic superconformal states.

Unfortunately, it is impossible to reproduce the character (4.117) using this assumption for the structure of κ and the assumption that the residues have a structure like the examples we have seen ((4.69) or (4.94)). Specifically, the latter assumption means that at level α of the filtration we form clusters of r_α particles; the wavefunction is taken to vanish as d_α powers as two clusters approach,

d'_α powers as one of the remaining particles approach a cluster, and d''_α powers as two remaining particles approach each other, for some sets of numbers d, d', d''. My result is that no choice of integers d, d', d'' reproduces the degree $\vec{m}A\vec{m}$ for a partition κ of this form.

I can obtain a solution if we relax some of these assumptions, for example if we let the widths of some blocks be wider than two or if we allow clusters in a block to interact with couplings of different degree with particles from different blocks. At this point, though, we clearly need to abandon the simple hypothesis set forth at the start of this section and reconsider how the filtration argument is to be set up in this situation.

Manual construction of the finite-N characters

However, we do not need to explicitly construct the basis states in order to count them; instead we may capitalize on the fact that the minimal SCFTs $SM(2,8)$ and $SM(3,5)$ (and only these two) are identical to the Virasoro minimal models $M(3,8)$ and $M(4,5)$. An overcomplete basis for the descendants of these states under the action of Virasoro and chiral boson modes is provided by the functional model first given in [158] — here I follow the specific presentation given in [68].

Particle operators in conformal block wavefunctions are given by the tensor product of a chiral vertex operator and the simple current of some rational CFT. Here I take this CFT to be a Virasoro minimal model (section 4.2.3) $M(p,p')$, whose only simple currents are the identity $\phi_{1,1}$ and $\phi_{1,p'-1} = \phi_{p-1,1}$ [159]. I choose the latter, since this operator is the supercurrent in both $M(3,8)$ and $M(4,5)$. The chiral vertex operator then needs to have charge β, with $\beta^2 = (p-2)(p'-2)/2$; a background charge of $\sqrt{2}(\beta - 1/\beta)$ is placed at infinity.

We saw before how the presence of singular vectors in a Verma module led to differential equations for amplitudes of the corresponding operator (equations (4.39), (4.40)). This action may be translated to the case of FQHE states built out of those operators by introducing

$$\ell_m = \sum_{j=1}^{N} z_j^{m+1} \frac{\partial}{\partial z_j}, \qquad m \geq -1, \tag{4.119}$$

$$p_m = \sum_{j=1}^{N} z_j^m, \qquad m \geq 0. \tag{4.120}$$

Here the ℓ_ms implement half of the Virasoro algebra (4.33), with $[\ell_m, \ell_n] = (m-n)\ell_{m+n}$, and the power sum symmetric polynomials arise from the chiral vertex operators.

For the model $SM(2,8) = M(3,8)$, the supercurrent is the $\phi_{2,1}$ Virasoro operator which has a null vector at level two given by (*e.g.*, [135])

$$\left[8L_{-2} - 3L_{-1}^2\right] |\phi_{2,1}\rangle = 0. \tag{4.121}$$

In terms of polynomial wavefunctions, when the contribution of the chiral boson is added this becomes the differential equation

$$\left[16\ell_2 - 6\ell_1^2 + 36(N-2)p_1\ell_1 - (54N^2 - 216N + 240)p_1^2 - 6(5N-18)p_2\right]\varphi_N = 0. \tag{4.122}$$

(this agrees with equation (3.6) of [68] for $p = 8, N = 3$). In the tricritical Ising model $SM(3,5) = M(4,5)$, the supercurrent is the Virasoro field $\phi_{3,1}$ which has the singular vector [135]

$$\left[15L_{-3} - 20L_{-2}L_{-1} + 4L_{-1}^3\right] |\phi_{3,1}\rangle = 0, \tag{4.123}$$

which translates into the following differential equation:

$$\begin{aligned}\big[&15\ell_3 - 20\ell_2\ell_1 + 4\ell_1^3 + 60(N-2)p_1\ell_2 - 36(N-2)p_1\ell_1^2 \\ &+ 6(18N^2 - 72N + 77)p_1^2\ell_1 + 12(2N-9)p_2\ell_1 - 18(N-2)(6N^2 - 24N + 29)p_1^3 \\ &-9(8N^2 - 52N + 77)p_2p_1 - 9(N-12)p_3\big]\varphi_N = 0.\end{aligned} \tag{4.124}$$

This implies that we have a linear relation between the ten basis states $(\ell_3\varphi_N, \ell_2\ell_1\varphi_N, \ldots, p_3\varphi_N)$ at level three. In fact, we already know that for small N there are *more* relations than that: the number of linearly-independent states is given by the character formula for our polynomial basis at finite N, equation (4.101) — note that although this counting was done for a different basis, both

span the same spaces with the same dimensions. The characters in the odd vacuum sector have the following expansions:

$$\text{ch } I_3^{\text{SCFT}} = 1 + 2q + 4q^2 + 6q^3 + 8q^4 + 10q^5 + 13q^6 + \cdots \tag{4.125}$$

$$\text{ch } I_5^{\text{SCFT}} = 1 + 2q + 5q^2 + 9q^3 + 16q^4 + 24q^5 + 36q^6 + \cdots \tag{4.126}$$

$$\text{ch } I_7^{\text{SCFT}} = 1 + 2q + 5q^2 + 10q^3 + 19q^4 + 32q^5 + 54q^6 + \cdots \tag{4.127}$$

$$\text{ch } I_{\infty,\,\text{odd}}^{\text{SCFT}} = 1 + 2q + 5q^2 + 10q^3 + 20q^4 + 36q^5 + 66q^6 + \cdots \tag{4.128}$$

We may compare these with the character for the tricritical Ising model, for example in the forms (4.116) or (4.117), after adding in the contribution of the chiral boson:

$$\frac{\chi_{3,1}^{(4,5)}}{(q)_\infty} = 1 + 2q + 5q^2 + 9q^3 + 17q^4 + 29q^5 + 50q^6 + \cdots \tag{4.129}$$

Inspection confirms that in order to obtain the tricritical Ising model, states must be removed from the Kac module starting at level three. The analogues of (4.128) and (4.129) for N even are

$$\text{ch } I_{\infty,\,\text{even}}^{\text{SCFT}} = 1 + q + 3q^2 + 5q^3 + 11q^4 + 18q^5 + 35q^6 + \cdots, \tag{4.130}$$

$$\frac{\chi_{1,1}^{(4,5)}}{(q)_\infty} = 1 + q + 3q^2 + 5q^3 + 10q^4 + 16q^5 + 29q^6 + \cdots. \tag{4.131}$$

In what follows, I will concentrate on the N odd, since a singular vector appears at lower level there (three versus four) and must be dealt with first.

It is unclear from the above considerations how many states need to be removed at small N: strictly speaking, equation (4.129) only applies in the $N \to \infty$ limit, when we can use an arbitrarily large number of modes to construct states in our basis. In the absence of an explicit basis, we may still use the functional model given above to compute the dimension of the irreducible module of the tricritical Ising model, starting with the amplitude φ_N, computing its descendants under the operators ℓ_m and p_m, and counting the number of linear relations between the resulting set of polynomials, level by level. In fact, we can do this for *all* SCFTs, since the state at level zero for

arbitrary c is given by Simon's result (4.86), so we may leave c as a free parameter.

We can only carry out this computation so far, though, because we are only working with Virasoro and chiral boson modes, and the character of the basis used in the functional model (all combinations of the p_ms and ℓ_ms acting on φ_N, modulo their mutual commutation relations) is

$$\text{ch } I_\infty^{\text{Vir}} = \frac{1}{(q)_\infty^2} = 1 + 2q + 5q^2 + 10q^3 + 20q^4 + 36q^5 + 65q^6 + \cdots . \tag{4.132}$$

At level six and higher there are more states in I_∞^{SCFT} than we can create in the functional model, so the calculation is only valid up to level five. In practice, going to high levels with even moderate numbers of particles proved prohibitive with the computer resources I had available.

Numerical results

With three particles, at any level I only obtained linear relations beyond those implied by (4.125) for values of A corresponding to the $M(3,8)$ model (and the $c = 0$ model, which is a Laughlin state without true three-body interactions and hence uninteresting for our purposes, as noted above). Since the finite-N character for this model has already been found in equation (4.111), I do not need to elaborate on the $N = 3$ case.

For $N = 5$ and 7, I simplify the calculations by working within the translationally-invariant subspace of the functional model's basis. I do this by constructing states from the following translationally invariant verisons of (4.119), (4.120):

$$\hat{p}_m(\{z_i\}) = p_m\Big(\{z_i - \frac{1}{N}\sum_{j=1}^{N} z_j\}\Big), \tag{4.133}$$

$$\hat{\ell}_m = \ell_m + \sum_{k=1}^{m} \frac{(-1)^k}{N}\binom{m+1}{k} p_k \ell_{m-k}. \tag{4.134}$$

Both of these vanish when acted on by $\ell_{-1} = \sum_{j=1}^{N} \partial_{z_j}$, the generator of translations in the plane.

We cannot encounter a singular vector in the finite-N module of the tricritical Ising model before level three. Here there are five translationally-invariant basis states in the functional model,

which I took to be

$$\hat{v}_1 = \hat{\ell}_3\varphi_N, \quad \hat{v}_2 = \hat{\ell}_2\hat{\ell}_1\varphi_N, \quad \hat{v}_3 = \hat{\ell}_1^3\varphi_N, \quad \hat{v}_4 = \hat{p}_2\hat{\ell}_1\varphi_N, \quad \hat{v}_5 = \hat{p}_3\varphi_N. \tag{4.135}$$

We may check for linear relations among these by simply evaluating them at five distinct sets of coordinates $z_1, \ldots, z_N$ and finding the zeros of the determinant of the resulting matrix. Performing this confirms that the singular vector relation (4.124) is redundant with the finite-size constraints at $N = 5$: we do not "see" the singular vector with this many particles. However, the extra linear relation is present at $N = 7$, as we can tell from comparing the characters (4.127) and (4.129). Therefore, we have the following results for the finite-N version of (4.129):

$$\begin{aligned}
(q)_3^{-1}\ \chi_{3,1}^{(4,5)}\Big|_{N=3} &= 1 + 2q + 4q^2 + 6q^3 + \cdots \\
(q)_5^{-1}\ \chi_{3,1}^{(4,5)}\Big|_{N=5} &= 1 + 2q + 5q^2 + 9q^3 + \cdots \\
(q)_7^{-1}\ \chi_{3,1}^{(4,5)}\Big|_{N=7} &= 1 + 2q + 5q^2 + 9q^3 + \cdots
\end{aligned} \tag{4.136}$$

The fact that the singular vector is only seen at $N = 7$ implies that the extra term in the projection Hamiltonian required to remove it from the generic superconformal basis will involve interactions among seven particles, presumably in some sort of clustered configuration. At the time of writing I have no way to further constrain what form this interaction must take.

I also performed the calculation for $N = 7$ by projecting the basis states onto translationally-invariant combinations of states in the generic superconformal basis in an attempt to shed light on the form of the states that must be removed at level three and higher. In that basis, I found that the singular vector at level three, $N = 7$ is given by

$$\begin{aligned}
V_{\text{sing}} = &-101680\Phi_{(3,0,1)}p_3^{(1)} + 51480\Phi_{(3,0,1)}p_{21}^{(1)} - 5760\Phi_{(3,0,1)}p_{111}^{(1)} \\
&- 3840\Phi_{(3,0,1)}p_2^{(1)}p_1^{(3)} + 720\Phi_{(3,0,1)}p_{11}^{(1)}p_1^{(3)} - 480\Phi_{(3,0,1)}p_1^{(1)}p_{11}^{(3)} + 960\Phi_{(3,0,1)}p_{111}^{(3)} \\
&+ 147\Phi_{(2,1,1)}p_1^{(1)} - 392\Phi_{(2,1,1)}p_1^{(2)} + 196\Phi_{(2,1,1)}p_1^{(3)}.
\end{aligned} \tag{4.137}$$

Here $p_\lambda^{(\alpha)}$ is the power-sum symmetric polynomial in the center-of-mass coordinates of the clusters

in block α of (4.97), and Ψ is $\Psi_\kappa^{\rm SCFT}$ evaluated at the value of A corresponding to the central charge of the tricritical Ising model. The partition κ is specified by the lengths $\{m_1, m_2, m_3\}$ of the blocks. Equation (4.137) must be a solution of the singular vector differential equation (4.124), but other than that no useful features are evident.

4.4 Discussion

As the reader doubtlessly already surmised, much work remains to be done on this topic. The partial results on the tricritical Ising model discussed in the previous section should imply that the structure of states for a given superconformal minimal model will look very different than those for the generic theory (4.97). The fact that seven-particle interactions are necessary to obtain the tricritical Ising SCFT hints strongly that some scenario involving clusters of clusters is taking place.

Constructing the finite-N characters and singular vectors by hand, as was done in the last section, may or may not be a productive enterprise. While we are limited in the maximum level of the module that we can construct, I could conceivably extend the characters (4.136) to higher N by formulating the calculation in terms of the algebra of symmetric functions, rather than representing each polynomial as an explicit function of coordinates. Any real advance will come from constructing a basis with the singular vector conditions built in in some way.

A situation intermediate in difficulty between $SM(2,8)$ and the tricritical Ising model may be provided by the $SM(2,4k)$ series of superconformal minimal models with $k \geq 2$, whose characters were shown to have a simple form without q-binomials [160]:

$$\hat{\chi}_{1,1}^{(2,4k)}(q) = \sum_{m_1,\ldots,m_{2k-2}} \frac{q^{\frac{1}{2}(M_1^2+M_2^2+\cdots+M_{2k-2}^2)+M_1+M_3+\cdots+M_{2k-3}}}{(q)_{m_1}(q)_{m_2}\cdots(q)_{m_{2k-2}}} \tag{4.138}$$

with $M_j = m_j + m_{j+1} + \cdots + m_{2k-2}$. The case $k = 2$ reproduces the character for $SM(2,8)$ given above (4.112). For $k = 3$ we appear to run into the same sorts of problems encountered in examining the tricritical ising character: namely, that no consistent set of integer powers d, d', d'' reproduces the degree of the polynomial implied by (4.138) for a partition κ with blocks of width two or one.

The simplicity of the characters (4.138) reflect a simplicity in the structure of the irreducible modules of these theories: in [161] it was shown how to construct all modules of this theory using modes of the supercurrent subject to simple exclusion conditions. This points to the fundamental matter underlying this project: there still lacks some understanding of how to extract the necessary information from the CFT in question. Clustering properties are obviously related to OPEs of the corresponding fields, but because of the chiral boson the OPEs of the particle operators are entirely nonsingular; we have to know this OPE to high order — exactly how high is unclear. This is how the constraint (2.5) is derived in [68], but their main argument makes use of different (filtration) methods. Nevertheless, in [159] an iterative procedure was obtained to construct all descendant operators appearing in the OPE of the simple current field $\phi_{1,p-1}$ with itself in any Virasoro minimal model, and this argument may be useful for the tricritical Ising case.

Bibliography

[1] T. S. Jackson and N. Read, Phys. Rev. E **81**, 021130 (2010), arXiv:0902.3651.

[2] T. S. Jackson and N. Read, Phys. Rev. E **81**, 021131 (2010), arXiv:0909.5343.

[3] E. Lawler, *Combinatorital Optimization: Networks and Matroids* (Dover Publications, Mineola, NY, 2001).

[4] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Dover Publications, Mineola, NY, 1998).

[5] R. E. Tarjan, *Data Structures and Network Algorithms* (Society for Industrial and Applied Mathematics, Philadeplphia, PA, 1983).

[6] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. S. Schrijver, *Combinatorial Optimization* (Wiley Interscience, New York, NY, 1998).

[7] J. M. Steele, *Probability Theory and Combinatorial Optimization* (Society for Industrial and Applied Mathematics, Philadeplphia, PA, 1997).

[8] J. B. Kruskal, Proc. Amer. Math. Soc. **7**, 48 (1956).

[9] V. Jarník, Práce Moravské Příirodovědecké Společnosti **6**, 57 (1930).
R. C. Prim, Bell System Tech. J. **36**, 1389 (1957).
E. W. Dijkstra, Numerische Mathematik **1**, 269 (1959).

[10] R. Dobrin and P. M. Duxbury, Phys. Rev. Lett. **86**, 5076 (2001).

[11] M. Cieplak, A. Maritan, and J. R. Banavar, Phys. Rev. Lett. **72**, 2320 (1994).
M. Cieplak, A. Maritan, and J. R. Banavar, Phys. Rev. Lett. **76**, 3754 (1996).
M. Cieplak, A. Maritan, M. R. Swift, A. Bhattacharya, A. L. Stella, and J. R. Banavar, J. Phys. A **28**, 5693 (1995).

[12] S. Sreenivasan, T. Kalisky, L. A. Braunstein, S. V. Buldyrev, S. Havlin, and H. E. Stanley, Phys. Rev. E **70**, 046133 (2004).

[13] G. J. Sazbó, M. Alava, and J. Kertész, Physica A **330**, 31 (2003).

[14] x LAN/MAN standards committee of the IEEE computer society, ed., *ANSI/IEEE Std. 802.1D* (IEEE, 1990).

[15] D. Aldous and J. Steele, Probab. Theory Relat. Fields **92**, 247 (1992).

[16] J. C. Gower and G. J. S. Ross, J. Appl. Stat. **18**, 54 (1969).
G. J. S. Ross, J. Appl. Stat. **18**, 103 (1969).
C. T. Zahn, IEEE Trans. Comput. C **20**, 68 (1971).

[17] A.-L. Barabási, Phys. Rev. Lett. **76**, 3750 (1996).

[18] A. A. Middleton, Phys. Rev. B **61**, 14787 (2000).

[19] E. López, S. V. Buldyrev, L. A. Braunstein, S. Havlin, and H. E. Stanley, Phys. Rev. E **72**, 056131 (2005).
S. Buldyrev, S. Havlin, and H. E. Stanley, Phys. Rev. E **73**, 036128 (2006).
M. Porto, N. Schwartz, S. Havlin, and A. Bunde, Phys. Rev. E **60**, 2448 (1999).

[20] Z. Wu, L. A. Braunstein, S. Havlin, and H. E. Stanley, Phys. Rev. Lett. **96**, 148702 (2006).
L. A. Braunstein, Z. Wu, Y. Chen, S. V. Buldyrev, S. Sreenivasan, T. Kalisky, R. Cohen, E. Lopez, S. Havlin, and H. E. Stanley, Int. J. Bifurcations and Chaos **17**, 2215 (2007), `cond-mat/0606338v4`.

[21] N. Read, Phys. Rev. E **72**, 036114 (2005).

[22] Z. Wu, E. Lopez, S. V. Buldyrev, L. A. Braunstein, S. Havlin, and H. E. Stanley, Phys. Rev. E **71**, 045101 (2005).

[23] C. M. Newman and D. L. Stein, Phys. Rev. Lett. **72**, 2286 (1994).
C. M. Newman and D. L. Stein, J. Stat. Phys. **82**, 1113 (1996).

[24] Y. Fu and P. W. Anderson, J. Phys. A **19**, 1605 (1986).

[25] M. Mézard, G. Parisi, and M. Virasoro, eds., *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).

[26] M. Mézard and A. Montanari, *Information, Physics and Computation* (Oxford University Press, Oxford, 2009).

[27] O. Bor**r**uvka, Práce Moravské Příirodovědecké Společnosti v Brně **3**, 37 (1926).

[28] R. L. Graham and P. Hell, IEEE Ann. Hist. Comput. **7**, 43 (1985).

[29] J. Eisner (1997), manuscript available online (78 pages), University of Pennsylvania, URL `http://cs.jhu.edu/~jason/papers/#ms97`.

[30] L. A. Braunstein, S. V. Buldyrev, R. Cohen, S. Havlin, and H. E. Stanley, Phys. Rev. Lett. **91**, 168701 (2003).

[31] S. V. Buldyrev, S. Havlin, E. López, and H. E. Stanley, Phys. Rev. E **70**, 035102(R) (2004).

[32] Y. M. Strelniker, R. Berkovits, A. Frydman, and S. Havlin, Phys. Rev. E **69**, 065105 (2004).

[33] B. Wieland and D. B. Wilson, Phys. Rev. E **68**, 056101 (2003).

[34] C. M. Newman and D. L. Stein, Phys. Rev. E **63**, 016101 (2000).

[35] A. A. Middleton, Phys. Rev. Lett. **83**, 1672 (1999).

[36] J. Houdayer and O. C. Martin, Europhys. Lett. **49**, 794 (2000).

[37] O. L. White and D. S. Fisher, Phys. Rev. Lett. **96**, 137204 (2006).

[38] O. Motrunich, S.-C. Mau, D. A. Huse, and D. S. Fisher, Phys. Rev. B **61**, 1160 (2000).

[39] D. Stauffer and A. Aharony, *Introduction to Percolation Theory* (Taylor and Francis, London, 1994), 2nd ed.

[40] J. W. Essam, Rep. Prog. Phys. **43**, 833 (1980).

[41] K. S. Alexander, Ann. Prob. **23**, 87 (1995).

[42] M. Aizenman, A. Burchard, C. M. Newman, and D. B. Wilson, Random Struct. Algorithms **15**, 319 (1999).
D. Wilson, Phys. Rev. E **69**, 037105 (2004).

[43] C. Garban, G. Pete, and O. Schramm, `arXiv:0909.3138`.

[44] D. Aldous, Probab. Theory Relat. Fields **93**, 507 (1992).

[45] D. Aldous, Random Struct. Algorithms **18**, 381 (2001).

[46] M. Aizenman, Nucl. Phys. B **485**, 551 (1997).

[47] R. Pemantle, Ann. Prob. **19**, 1559 (1991).

[48] R. Lyons, Y. Peres, and O. Schramm, Ann. Prob. **34**, 1665 (2006).

[49] T. Hara and G. Slade, Commun. Math. Phys. **128**, 333 (1990).
G. Slade, *The Lace Expansion and its Applications* (Springer, New York, NY, 2004).

[50] R. B. Laughlin, Phys. Rev. Lett. **50**, 1395 (1983).

[51] S. H. Simon, E. H. Rezayi, and N. R. Cooper, Phys. Rev. B **75**, 075318 (2007).

[52] S. H. Simon, J. Phys. A **42**, 055402 (2009).

[53] N. Read, Phys. Rev. B **79**, 245304 (2009).

[54] X.-G. Wen and Y.-S. Wu, Nucl. Phys. B **419**, 455 (1994).
X.-G. Wen, Y.-S. Wu, and Y. Hatsugai, Nucl. Phys. B **422**, 476 (1994).

[55] X.-G. Wen and Z. Wang, Phys. Rev. B **77**, 235108 (2008).
X.-G. Wen and Z. Wang, Phys. Rev. B **78**, 155109 (2008).
M. Barkeshli and X.-G. Wen, Phys. Rev. B **79**, 195132 (2009).
Y.-M. Lu, X.-G. Wen, Z. Wang, and Z. Wang, `arXiv:0910.3988`.

[56] I. G. Macdonald, *Symmetric functions and Hall polynomials* (Clarendon Press, Oxford, 1998), 2nd ed.

[57] B. L. Feigin, M. Jimbo, T. Miwa, and E. Mukhin, Int. Math. Res. Not. **23**, 1223 (2002).

[58] B. A. Bernevig and F. D. M. Haldane, Phys. Rev. Lett. **100**, 246802 (2008).
B. A. Bernevig and F. D. M. Haldane, Phys. Rev. B **77**, 184502 (2008).
B. A. Bernevig and F. D. M. Haldane, Phys. Rev. Lett. **102**, 066802 (2009).

[59] E. J. Bergholtz, T. H. Hansson, M. Hermanns, and A. Karlhede, Phys. Rev. Lett. **99**, 256803 (2007).
E. J. Bergholtz, T. H. Hansson, M. Hermanns, A. Karlhede, and S. Viefers, Phys. Rev. B **77**, 165325 (2008).
E. Ardonne, E. J. Bergholtz, J. Kailasvuori, and E. Wikberg, J. Stat. Mech. **0804**, P04016 (2008).
E. Ardonne, Phys. Rev. Lett. **102**, 180401 (2009).

[60] B. A. Bernevig, V. Gurarie, and S. H. Simon, J. Phys. A **42**, 245206 (2009).
B. Estienne, N. Regnault, and R. Santachiara, Nucl. Phys. B **824**, 539 (2010).

[61] B. Estienne and R. Santachiara, J. Phys. A **42**, 445209 (2009).

[62] N. Read and E. H. Rezayi, Phys. Rev. B **59**, 8084 (1999).

[63] S. H. Simon, E. H. Rezayi, and N. Regnault, `arXiv:0908.0947`.

[64] S. H. Simon, E. H. Rezayi, and N. R. Cooper, Phys. Rev. B **75**, 195306 (2007).

[65] N. Read, Phys. Rev. B **73**, 245334 (2006).

[66] N. Read and E. H. Rezayi, Phys. Rev. B **54**, 16864 (1996).

[67] N. Read (2009), *Notes on zero-energy states for three-body special Hamiltonian forbidding degree less than, and one term at, degree six* (unpublished).

[68] B. L. Feigin, M. Jimbo, and T. Miwa, Prog. Math. Phys. **23**, 179 (2002).

[69] Z. Qiu, Nucl. Phys. B **270**, 205 (1986).

[70] B. Chazelle, J. ACM **47**, 1028 (2000).

[71] N. Read, Phys. Rev. E **70**, 027103 (2004).

[72] G. Grimmett, *Percolation* (Springer, New York, NY, 1999), 2nd ed.

[73] R. Lenormand and S. Bories, C. R. Acad. Sci. **291**, 279 (1980).
R. Chandler, J. Koplick, K. Lerman, and J. F. Willemsen, J. Fluid Mech. **119**, 249 (1982).
D. Wilkinson and J. Willemsen, J. Phys. A **16**, 3365 (1983).

[74] J. T. Chayes, L. Chayes, and C. M. Newman, Commun. Math. Phys. **101**, 383 (1985).

[75] A. M. Frieze, Discrete App. Math. **10**, 47 (1985).

[76] K. H. Fischer and J. A. Hertz, *Spin Glasses* (Cambridge University Press, Cambridge, 1993).

[77] C. M. Newman and D. L. Stein, J. Phys: Cond. Matt. **15**, R1319 (2003).

[78] S. Edwards and P. W. Anderson, J. Phys. F **5**, 965 (1975).

[79] F. Barahona, J. Phys. A **15**, 3241 (1982).

[80] C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1984).

[81] D. Sherrington and S. Kirkpatrick, Phys. Rev. Lett. **35**, 1792 (1975).

[82] D. J. Thouless, P. W. Anderson, and R. J. Palmer, Phil. Mag. **35**, 593 (1977).

[83] G. Parisi, Phys. Rev. Lett. **43**, 1754 (1979).
G. Parisi, Phys. Rev. Lett. **50**, 1946 (1983).
M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, Phys. Rev. Lett. **52**, 1156 (1984).
M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, J. Phys (Paris) **45**, 843 (1984).

[84] A. J. Bray and M. A. Moore, Phys. Rev. B **31**, 631 (1985).
A. J. Bray and M. A. Moore, Phys. Rev. Lett. **58**, 57 (1987).
D. S. Fisher and D. A. Huse, Phys. Rev. Lett. **56**, 1601 (1986).
D. A. Huse and D. S. Fisher, J. Phys. A **20**, L997 (1987).
D. S. Fisher and D. A. Huse, J. Phys. A **20**, L1005 (1987).
D. S. Fisher and D. A. Huse, Phys. Rev. B **38**, 386 (1988).

[85] M. E. Fisher and J. W. Essam, J. Math. Phys. **2**, 609 (1961).

[86] K. B. Arthreya and P. E. Ney, *Branching Processes* (Addison-Wesley, Springer-Verlag, Berlin, 1972; Dover Publications, Mineola, NY, 2004).

[87] M. Steele and J. A. Fill, in *Stein's Method and Applications*, edited by A. Barbour and L. Chen (World Publications, Singapore, 2005).

[88] F. Y. Wu, Rev. Mod. Phys. **54**, 235 (1982).

[89] R. K. P. Zia and D. J. Wallace, J. Phys. A **8**, 1495 (1975).

[90] T. C. Lubensky, in *Ill-Condensed Matter*, edited by R. Balian, R. Maynard, and G. Toulouse (North-Holland, New York, 1979), Les Houches session XXXI.

[91] W. T. Tutte, J. Combinatorial Theory **2**, 301 (1967).
W. T. Tutte, J. Combinatorial Theory **9**, 289 (1970).

[92] J. W. Essam, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and M. S. Green (Academic Press, New York, 1972), vol. 2.

[93] J. H. van Lint and R. M. Wilson, *A Course in Combinatorics* (Cambridge U. Press, Cambridge, 2001), 2nd ed.

[94] A. Coniglio and J. W. Essam, J. Phys. A **10**, 1917 (1977).

[95] D. K. Arrowsmith, J. Math. Phys. **20**, 101 (1979).

[96] D. J. Amit, J. Phys. A **9**, 1441 (1976).

[97] R. G. Priest and T. C. Lubensky, Phys. Rev. B **13**, 4159 (1976).
R. G. Priest and T. C. Lubensky, Phys. Rev. B **14**, 5125 (1976).

[98] P. W. Kasteleyn and C. M. Fortuin, J. Phys. Soc. Japan Suppl. **26**, 11 (1969).
C. M. Fortuin and P. W. Kasteleyn, Physica **57**, 536 (1972).

[99] O. F. de Alcantara Bonfim, J. E. Kirkham, and A. J. McKane, J. Phys. A **14**, 2391 (1981).

[100] C. Itzykson and J.-M. Drouffe, *Statistical Field Theory*, vol. 1 (Cambridge University Press, Cambridge, 1989).

[101] C. Itzykson and J.-B. Zuber, *Quantum Field Theory* (McGraw-Hill, New York, NY, 1980).

[102] D. Amit and V. Martín-Mayor, *Field Theory, the Renormalization Group, and Critical Phenomena* (World Scientific, Hackensack, NJ, 2005), 3rd ed.

[103] A. Coniglio, J. Phys. A **15**, 3829 (1982).

[104] H. E. Stanley, J. Phys. A **10**, L211 (1977).

[105] J. Komlós, Combinatorica **5**, 57 (1985).
B. Dixon, M. Rauch, and R. E. Tarjan, SIAM J. Comp. **21**, 1184 (1992).
V. King, Algorithmica **18**, 263 (1997).

[106] D. R. Karger, N. Klein, and R. E. Tarjan, J. ACM **42**, 321 (1995).
M. Bergère and Y. M. P. Lam, J. Math. Phys. **17**, 1546 (1976).

[107] N. Nakanishi, *Graph Theory and Feynman Integrals* (Gordon and Breach, New York, NY, 1970).

[108] T. Appelquist, Ann. Phys. **54**, 27 (1969).
M. Bergère and J. B. Zuber, Comm. Math. Phys. **35**, 113 (1974).

[109] N. Bogoliubov and O. Parasiuk, Acta Math. **97**, 227 (1957).

[110] K. Hepp, Comm. Math. Phys. **2**, 301 (1966).

[111] W. Zimmermann, Ann. Phys. **77**, 536 (1973).

[112] J. C. G. Callan, in *Methods in Field Theory*, edited by R. Balian and J. Zinn-Justin (North-Holland, Amsterdam, 1976), Les Houches lectures, Session XXVIII, p. 41.

[113] J. Illiopoulos, C. Itzykson, and A. Martin, Rev. Mod. Phys. **47**, 165 (1975).

[114] E. Brézin, J. C. le Gillou, and J. Zinn-Justin, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and M. S. Green (Academic Press, New York, 1976), vol. 6.

[115] J. L. Cardy and P. Grassberger, J. Phys. A **18**, L267 (1985).
K. H. Janssen, Z. Phys. B **58**, 311 (1985).
A. B. Harris, Phys. Rev. B **35**, 5056 (1987).
R. Blumenfeld, Y. Meir, A. B. Harris, and A. Aharony, J. Phys. A **19**, L791 (1986).
Y. Meir and A. B. Harris, Phys. Rev. Lett. **63**, 2819 (1989).

[116] T. Kalisky, S. Sreenivasan, L. A. Braunstein, S. V. Buldyrev, S. Havlin, and H. E. Stanley, Phys. Rev. E **73**, 025103(R) (2006).

[117] C. M. Bender and S. A. Orszag, *Advanced mathematical methods for scientists and engineers* (Springer, Berlin, 1999).

[118] A. Houghton, J. S. Reeve, and D. J. Wallace, Phys. Rev. B **17**, 2956 (1978).

[119] G. Moore and N. Read, Nucl. Phys. B **360**, 362 (1991).

[120] R. E. Prange and S. M. Girvin, eds., *The Quantum Hall Effect* (Springer, Berlin, 1990), 2nd ed.
S. M. Girvin, in *Topological aspects of low dimensional systems*, edited by A. Comtet, T. Jolicoeur, S. Ouvry, and F. David (Springer, Berlin, 2000).
M. Stone, ed., *Quantum Hall Effect* (World Scientific, Singapore, 1992).

[121] E. Fradkin, *Field theories of condensed matter systems*, vol. 82 of *Frontiers in Physics* (Westview, New York, 1991).
X.-G. Wen, *Quantum field theory of many-body systems* (Oxford University Press, Oxford, 2004).

[122] Q. Niu, D. J. Thouless, and Y.-S. Wu, Phys. Rev. B **31**, 3372 (1985).

[123] B. I. Halperin, Phys. Rev. B **25**, 2185 (1982).

[124] F. D. M. Haldane, Phys. Rev. Lett. **51**, 605 (1983).
F. D. M. Haldane, in *The Quantum Hall Effect*, edited by R. E. Prange and S. M. Girvin (Springer, New York, 1990).

[125] S. M. Girvin and T. Jach, Phys. Rev. B **29**, 5617 (1984).

[126] N. R. Cooper, N. K. Wilkin, and J. M. F. Gunn, Phys. Rev. Lett. **87**, 120405 (2001).
E. H. Rezayi, N. Read, and N. R. Cooper, Phys. Rev. Lett. **95**, 160404 (2005).
N. R. Cooper, Adv. Phys. **57**, 539 (2008).

[127] F. D. M. Haldane, Phys. Rev. Lett. **51**, 605 (1983).
B. I. Halperin, Phys. Rev. Lett. **52**, 1583 (1984).

[128] D. Arovas, J. R. Schrieffer, and F. Wilczek, Phys. Rev. Lett. **53**, 722 (1984).

[129] M. Milovanović and N. Read, Phys. Rev. B **53**, 13559 (1996).

[130] S. M. Girvin and A. H. MacDonald, Phys. Rev. Lett. **58**, 1252 (1987).

[131] S. C. Zhang, T. H. Hansson, and S. Kivelson, Phys. Rev. Lett. **62**, 82 (1989).

[132] N. Read, Phys. Rev. Lett. **62**, 86 (1989).

[133] N. Read, Phys. Rev. B **79**, 045308 (2009).
N. Read, `arXiv:0807.3107`.

[134] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. D. Sarma, Rev. Mod. Phys. **80**, 1083 (2008).

[135] P. D. Francesco, P. Mathieu, and D. Sénéchal, *Conformal field theory* (Springer, New York, NY, 1997).

[136] J. L. Cardy, in *Fields, strings and critical phenomena*, edited by E. Brézin and J. Zinn-Justin (North-Holland, Amsterdam, 1990).
P. Ginsparg, in *Fields, strings and critical phenomena*, edited by E. Brézin and J. Zinn-Justin (North-Holland, Amsterdam, 1990).

[137] A. A. Belavin, A. M. Polyakov, and A. B. Zamolodchikov, Nucl. Phys. B **241**, 333 (1984).

[138] M. A. Virasoro, Phys. Rev. D **1**, 2933 (1970).

[139] V. G. Kac, in *Lecture Notes in Physics v. 94* (Springer, Berlin, 1979).

[140] B. L. Feigin and D. B. Fuchs, Funct. Anal. and Appl. **16**, 114 (1982).
B. L. Feigin and D. B. Fuchs, Funct. Anal. and Appl. **17**, 241 (1983).

[141] D. Friedan, Z. Qiu, and S. Shenker, Phys. Rev. Lett. **52**, 1575 (1984).
D. Friedan, Z. Qiu, and S. Shenker, Commun. Math. Phys. **107**, 535 (1986).

[142] P. Ginsparg, Nucl. Phys. B **295**, 153 (1988).

[143] A. Rocha-Caridi, in *Vertex Operators in Mathematics and Physics*, edited by J. Lepowsky, S. Mandelstam, and I. M. Singer (Springer, Berlin, 1985).

[144] G. Moore and N. Seiberg, Phys. Lett. B **212**, 451 (1988).
G. Moore and N. Seiberg, Commun. Math. Phys. **123**, 177 (1989).

[145] G. Moore and N. Seiberg, Phys. Lett. B **220**, 422 (1989).

[146] E. Witten, Commun. Math. Phys. **121**, 351 (1989).

[147] J. M. Leinaas and J. Myrheim, Nuovo Cimento B **37**, 1 (1977).
F. Wilczek, Phys. Rev. Lett. **48**, 1144 (1982).

[148] J. Wess and B. Zumino, Phys. Lett. B **37**, 95 (1971).
E. Witten, Nucl. Phys. B **223**, 422 (1983).

[149] P. Goddard, A. Kent, and D. Olive, Commun. Math. Phys. **103**, 105 (1986).

[150] H. Eichenherr, Phys. Lett. B **151**, 26 (1985).
M. A. Bershadsky, V. G. Knizhnik, and M. G. Teitelman, Phys. Lett. B **151**, 31 (1985).
D. Friedan, Z. Qiu, and S. Shenker, Phys. Lett. B **151**, 37 (1985).

[151] R. P. Stanley, *Enumerative Combinatorics*, vol. 2 of *Cambridge studies in advanced mathematics 62* (Cambridge U. Press, Cambridge, 1999).

[152] E. Ardonne, R. Kedem, and M. Stone, J. Phys. A **38**, 617 (2005).

[153] A. Cappelli, L. S. Georgiev, and I. T. Todorov, Nucl. Phys. B **599**, 499 (2001).

[154] S. H. Simon, E. H. Rezayi, N. R. Cooper, and I. Berdnikov, Phys. Rev. B **75**, 075317 (2007).

[155] A. Berkovich, B. M. McCoy, and A. Schilling, Commun. Math. Phys. **191**, 325 (1998).

[156] E. Baver and D. Gepner, Phys. Lett. B **372**, 231 (1996).

[157] R. Kedem, T. R. Klassen, B. M. McCoy, and E. Melzer, Phys. Lett. B **304**, 263 (1993).

[158] B. L. Feigin and A. V. Stoyanovsky, Funct. Anal. and Appl. **28**, 55 (1993).

[159] P. Mathieu and D. Ridout, Nucl. Phys. B **776**, 365 (2007).

[160] E. Melzer, `hep-th/9412154`.
A. Berkovich, B. M. McCoy, and W. P. Orrick, J. Stat. Phys. **83**, 795 (1996).

[161] J.-F. Fortin, P. Jacob, and P. Mathieu, J. Phys. A **38**, 1699 (2005).