

**MARKOV CHAINS AND EMERGENT BEHAVIOR FOR PROBLEMS FROM  
DISCRETE GEOMETRY**

A Dissertation  
Presented to  
The Academic Faculty

By

Sarah Cannon

In Partial Fulfillment  
of the Requirements for the Degree  
Doctor of Philosophy in  
Algorithms, Combinatorics, and Optimization

School of Computer Science  
Georgia Institute of Technology  
August 2018

Copyright © Sarah Cannon 2018

## ACKNOWLEDGEMENTS

First, I would like to acknowledge the support and encouragement of my advisor, Dana Randall, throughout my time at Georgia Tech. Her help has been invaluable in navigating all aspects of grad school from start to finish.

I would also like to thank the other coauthors of the research presented in this thesis: Marta Andrés Arroyo, Joshua J. Daymude, David Levin, Sarah Miracle, Dana Randall, Andréa Richa, and Alexandre Stauffer. I especially thank Sarah Miracle, for helping me get started on research in this area, and Josh Daymude and Andrea Richa, for their close collaboration and willingness to explore new ideas.

I'd like to thank Sebastian Pokutta, Andréa Richa, Prasad Tetali, and Eric Vigoda for agreeing to be on my thesis committee, especially Eric Vigoda for being my thesis reader. I also appreciate Andréa, Prasad, and Eric's time and effort as my letter writers over the past year, and their help with making decisions about my future.

I would not have made it through graduate school without help from my friends, including the other students in Georgia Tech's Theory Lab and the members of the Georgia Tech Field Hockey Club and the Georgia Field Hockey Association.

Finally, I would like to acknowledge the unwavering support of my family: my parents, Sandy and Andy, and my sister Emily. Their support and encouragement from the beginning is the reason I am where I am today.

## TABLE OF CONTENTS

<b>Acknowledgments</b> . . . . .	iii
<b>List of Tables</b> . . . . .	ix
<b>List of Figures</b> . . . . .	x
<b>Chapter 1: Introduction</b> . . . . .	1
1.1 Random Sampling using Markov Chains . . . . .	2
1.1.1 Card Shuffling: Global and Local Moves . . . . .	3
1.1.2 The Hardcore Model from Statistical Physics . . . . .	4
1.2 Convergence Times of Markov Chains . . . . .	5
1.3 Biased Markov Chains . . . . .	6
1.3.1 Biased Card Shuffling . . . . .	7
1.3.2 The Weighted Hardcore Model from Statistical Physics . . . . .	7
1.4 Phase Transitions and Emergent Behavior . . . . .	8
1.5 Programmable Matter . . . . .	10
1.5.1 Self-organizing Particle Systems . . . . .	11
1.5.2 The Stochastic Approach . . . . .	11
1.5.3 Inspiration: Biological Systems . . . . .	12

<b>Chapter 2: Background</b>	14
2.1 Convergence to Stationarity: Mixing Time, Spectral Gap and Relaxation Time	15
2.2 Techniques for Proving Mixing and Relaxation Time Upper Bounds	17
2.2.1 Coupling	18
2.2.2 Path Coupling	18
2.2.3 Multiplicative Path Coupling	19
2.2.4 Comparison	20
2.2.5 Other Techniques	21
2.3 Techniques for Proving Mixing and Relaxation Time Lower Bounds	22
2.3.1 Diameter	22
2.3.2 Isoperimetric Inequalities	22
2.3.3 Distinguishing Statistics	23
<b>Chapter 3: Dyadic Tilings</b>	24
3.1 Preliminaries and Previous Work	26
3.1.1 Combinatorics of Dyadic Tilings	27
3.1.2 Markov Chains on Dyadic Tilings	30
3.2 The Edge-flip Markov Chain on Weighted Dyadic Tilings	31
3.3 Polynomial Convergence when $\lambda < 1$	34
3.4 Exponential Convergence when $\lambda > 1$	46
3.5 Polynomial Convergence when $\lambda = 1$	49
3.5.1 Proof Ideas	49
3.5.2 The Transition Matrix of $\mathcal{M}_k^{edge}$	51

3.5.3	The Block Dynamics Markov Chain $\mathcal{M}_k^{block}$ . . . . .	51
3.5.4	A Polynomial Upper Bound on the Mixing and Relaxation Times of $\mathcal{M}_k^{edge}$ . . . . .	59
3.5.5	A Nontrivial Lower Bound on Convergence when $\lambda = 1$ . . . . .	63
3.6	Weighted Dyadic Tilings and Statistical Physics . . . . .	70
3.7	Rectangular Dissections . . . . .	71
<b>Chapter 4: Free Boundary Planar Lattice Problems . . . . .</b>		<b>73</b>
4.1	Background on Grid 3-colorings . . . . .	74
4.1.1	Related Work: Sampling $k$ -colorings . . . . .	74
4.1.2	Sampling Grid 3-Colorings . . . . .	75
4.1.3	Problem Statement and Assumptions . . . . .	76
4.1.4	Glauber Dynamics and Tower Moves . . . . .	78
4.1.5	Height Functions . . . . .	79
4.2	A Markov Chain for Grid 3-colorings using Randomized Extensions . . . . .	82
4.2.1	Techniques . . . . .	83
4.2.2	A Sampling Algorithm using Randomized Extensions . . . . .	84
4.2.3	Extending One Boundary Segment . . . . .	85
4.2.4	Where Two Boundary Segments Meet: Convex Corner . . . . .	93
4.2.5	Where Two Boundary Segments Meet: Reflex Corner . . . . .	99
4.2.6	Properties of Markov Chain $\mathcal{M}_C$ . . . . .	101
4.2.7	Mixing Time of Markov Chain $\mathcal{M}_C$ . . . . .	104
4.2.8	Self-reducibility and Approximate Counting with $\mathcal{M}_C$ . . . . .	108
4.2.9	Using Approximate Counts to Sample from More Regions. . . . .	109

4.3	Lozenge Tilings . . . . .	110
<b>Chapter 5: Markov Chain Algorithms for Programmable Matter . . . . .</b>		<b>113</b>
5.1	The Stochastic Approach to Self-Organizing Particle Systems . . . . .	113
5.2	Background on Self-Organizing Particle Systems . . . . .	116
5.2.1	The Geometric Amoebot Model . . . . .	116
5.2.2	Terminology for Particle Systems . . . . .	118
5.2.3	Related Work: Particle Exclusion Processes . . . . .	120
5.3	Compression: Overview, Problem Definition, and Preliminaries . . . . .	120
5.3.1	Related Work . . . . .	122
5.3.2	Formalizing Compression: Perimeter and Edges . . . . .	123
5.4	Algorithms for Compression . . . . .	127
5.4.1	The Markov Chain $\mathcal{M}$ . . . . .	128
5.4.2	The Local Algorithm $\mathcal{A}$ . . . . .	129
5.4.3	Obliviousness and Robustness of $\mathcal{M}$ and $\mathcal{A}$ . . . . .	133
5.4.4	Invariants for Markov Chain $\mathcal{M}$ . . . . .	134
5.4.5	Eventual Ergodicity of Markov chain $\mathcal{M}$ . . . . .	137
5.4.6	The Stationary Distribution $\pi$ of $\mathcal{M}$ . . . . .	146
5.4.7	Convergence Time of Markov Chain $\mathcal{M}$ . . . . .	149
5.5	Achieving Compression and Expansion . . . . .	150
5.5.1	Preliminaries: Counting Particle Configurations by Perimeter . . . . .	151
5.5.2	Achieving Compression . . . . .	160
5.5.3	Using $\mathcal{M}$ for Expansion . . . . .	163

5.6	Shortcut Bridging: Overview, Problem Definition, and Preliminaries . . . .	168
5.6.1	Related Work . . . . .	169
5.6.2	Problem Definition . . . . .	170
5.6.3	The Stochastic Approach . . . . .	172
5.7	Algorithms for Shortcut Bridging . . . . .	174
5.7.1	Properties of Markov Chain $\mathcal{M}$ . . . . .	177
5.7.2	Simulations . . . . .	181
5.8	Dependence of Bridge Structure On Gap Angle . . . . .	183
5.8.1	Proofs for Small $\theta$ . . . . .	186
5.8.2	Proofs for Large $\theta$ . . . . .	191
5.9	Further Applications of the Stochastic Approach to Programmable Matter .	197
	<b>References . . . . .</b>	<b>199</b>

## LIST OF TABLES

- 4.1 Types of moves for regions with some free boundaries, and the probability with which they occur in  $\mathcal{M}_C$  and, specifically for rectangular regions, in  $\widetilde{\mathcal{M}}$  from [61]. Here  $q = \frac{1}{3|R' |}$  and  $s = \frac{1}{12mn}$ . . . . . 101



## LIST OF FIGURES

3.1	(a) A dyadic tiling of size 16 with a vertical bisector. (b) A dyadic tiling of size 16 with both a vertical and horizontal bisector. (c) A tiling that is not dyadic; the shaded rectangles are not dyadic because of their vertical dimension. . . . .	25
3.2	(a) A valid edge flip from one dyadic tiling to another. (b) Flipping the bold edge is not valid as the resulting tiling is not dyadic. (c) Flipping the bold edge is not valid as it does not result in a rectangular tiling. . . . .	25
3.3	State of Markov chain $\mathcal{M}_k^{edge}$ with $k = 6$ and $n = 64$ and various values of $\lambda$ after 1,000,000 simulated steps, starting from a tiling with all $1/n \times 1$ vertical rectangles. . . . .	34
3.4	A sequence of edge-flips from tiling $\sigma_h$ consisting of all $1 \times 1/n$ horizontal rectangles to ground state tiling $\sigma^*$ consisting of all $(1/\sqrt{n} \times 1/\sqrt{n})$ rectangles. . . . .	36
3.5	Rectangle $S$ of area $2/n$ in marginal tilings (a) $A_t$ and (b) $B_t$ . . . . .	38
3.6	An area $2/n$ rectangle $S$ bisected by (a) horizontal edge $e$ in $A_t$ and (b) vertical edge $f$ in $B_t$ . Four “bad” edge-flips $g, h, i, j$ exist only if $A_t$ and $B_t$ are tiled (up to reflection) in the neighborhood of $S$ as shown. . . . .	41
3.7	Possible configurations for the half-bisectors of $x$ and $y$ in Case 1 of the proof of Theorem 3.18. In figure (c), $i \in \{0, 1\}$ denotes how many grey quadrants are tiled identically in $x$ and $y$ . . . . .	55
3.8	Possible configurations for the half-bisectors of $x$ and $y$ in Case 2 of the proof of Theorem 3.18. The value of $i \in \{0, 1, 2, 3\}$ denotes the number of grey quadrants which are tiled identically in $x$ and $y$ . . . . .	57
3.9	Possible configurations for the half-bisectors of $x$ and $y$ in Case 3 of the proof of Theorem 3.18. The value of $i \in \{0, 1, 2, 3\}$ denotes the number of grey quadrants which are tiled identically in $x$ and $y$ . . . . .	57

3.10	The construction of a tiling to count $\prod_{i=0}^{k-2}  \Omega_i ^2$ . A rectangle with number $a$ indicates that we tile it with a tiling from $\Omega_{k-a}$ . . . . .	65
3.11	A tiling in $\partial\Omega_k^\downarrow$ , with the red edge being the flip that brings the tiling into $\Omega_k^\downarrow$ . . . . .	69
3.12	The construction of a tiling in $\partial\Omega_k^\downarrow$ . The grey areas represent the part that contains the pivotal edge. . . . .	70
4.1	(a) A sample input: a mixed boundary region $R$ for which we want to efficiently generate a uniformly random 3-coloring, and (b) one such proper 3-coloring of $R$ . . . . .	77
4.2	A 3-coloring of a rectangle $R$ and three tower moves, each of height 3. The top right tower abuts the boundary, the bottom tower is adjacent to the boundary, and the top left tower does both. The arrow drawn in each tower goes from the start cell to the end cell. . . . .	79
4.3	(a) a valid 3-coloring of a $7 \times 5$ rectangle, and (b) the height function obtained from it when fixing the height of the lowest leftmost cell to be its color. . . . .	80
4.4	(a) Three regions with height consistent mixed boundary conditions, where gray cells have fixed colors. (b) A region whose mixed boundary conditions are not height consistent; in the two colorings shown, the fixed cells have different heights. . . . .	81
4.5	(a) A portion $\ell$ of the boundary of a region $R$ lined by free boundary cells (with no fixed colors). (b) To obtain $R'$ we extend $R$ three units past $\ell$ . We will fix the colors adjacent to the boundary of this new region $R'$ . (c) $R$ 's up-up-down (UUD) coloring of the extended region $R'$ . (d) $R$ 's down-down-up (DDU) coloring of the extended region $R'$ . Throughout gray cells have fixed color. . . . .	86
4.6	Colorings (a) $\sigma$ and (b) $\tau$ differ by a tower of height 3 (grey) abutting boundary segment $\ell$ . (c) The indicated tower moves for the random colorings of the extended region $R'$ transition between $\sigma$ and $\tau$ . . . . .	89
4.7	Colorings (a) $\sigma$ and (b) $\tau$ differ at a single cell (grey) next to a boundary segment $\ell$ . (c) The indicated tower moves (dark gray, and indicated by arrows) and flip moves (indicated by circles) for the random colorings of the extended region $R'$ transition between $\sigma$ and $\tau$ . . . . .	91

4.8	Colorings (a) $\sigma$ and (b) $\tau$ differ by a tower of height 3 (grey) adjacent to a boundary segment $\ell$ . (c) The indicated tower moves for the random colorings of the extended region $R'$ transition between $\sigma$ and $\tau$ . . . . .	93
4.9	(a) A free boundary 3-coloring and the sides of one of its 16 extensions: the top and right sides are extended UUD, while the left and bottom sides are extended DDU. (b) The corners of the extension $R'$ and the deterministic colors given to them based on the random extension chosen. . . . .	94
4.10	(a) A region $R$ with a coloring $\sigma$ that has a convex corner and the region $R'$ we extend to near this convex corner. The four possible colorings of this extension in the neighborhood of this corner are shown in (b,c,d,e). Moves changing $\sigma$ into $\tau$ , which is the same as $\sigma$ except for its corner color, are shown in dark gray; tower moves are indicated by arrows, and flips are circled. . . . .	96
4.11	(a) A region $R$ with a coloring $\sigma$ that has a tower move including a convex corner of $R$ , and the region $R'$ we extend to near this convex corner. The four possible colorings of this extension in the neighborhood of this corner are shown in (b,c,d,e). Moves changing $\sigma$ into $\tau$ (Figure 4.12a) are shown in dark gray and indicated by arrows. . . . .	97
4.12	(a) A region $R$ with a coloring $\tau$ that has a tower move including a convex corner of $R$ , and the region $R'$ we extend to near this convex corner. The four possible colorings of this extension in the neighborhood of this corner are shown in (b,c,d,e). Moves changing $\tau$ into $\sigma$ (Figure 4.11a) are shown in dark gray and indicated by arrows. . . . .	98
4.13	(a) Two boundary segments of $R$ meeting a reflex corner, where the reflex cell (dark grey) has fixed color 0. (b) The extension of the vertical boundary segment, which includes extending the reflex cell. (c) The extension of the horizontal boundary segment, which includes extending the reflex cell. The extensions in (b) and (c) overlap, but are locally planar in the neighborhood of any free color cells. . . . .	100
4.14	(a) An intermediate step of the reduction between sampling and approximate counting, where the colors of some cells of a free boundary rectangle have been fixed. (b) The resulting mixed boundary region we sample from in one step in the approximate counting process. . . . .	109
4.15	(a) A free boundary lozenge tiling of a triangular region $E$ , and (b - c) the two possible random extensions for its top-right side. . . . .	111
4.16	A random extension, including all sides and completed corners, of a free boundary lozenge tiling of an equilateral triangle $E$ . . . . .	112

5.1	(a) A section of the triangular lattice $\Gamma$ ; (b) expanded particles (each denoted by its two occupied adjacent locations in $\Gamma$ and a thick line in between) and contracted particles (occupying one location); (c) two non-neighboring contracted particles with different offsets for their labels of the lattice edges leading to neighboring locations. . . . .	116
5.2	A particle configuration for which all valid moves of Markov chain $\mathcal{M}$ satisfy Property 2; no particle has a valid move satisfying Property 1. This demonstrates the subtlety of the Markov chain rules we have defined. . . . .	138
5.3	(a) An example of a particle configuration and a line $m_i$ that satisfies both invariants. (b) After a sequence of moves described in Lemma 5.9, $m_{i+1}$ satisfies Invariant 1. (c) After a sequence of moves described in Lemma 5.10, $m_{i+1}$ also satisfies Invariant 2. . . . .	139
5.4	Particle positions from the base case (top row) and inductive step (bottom row) of the proof of Lemma 5.8. Particles are represented by black circles, and unoccupied locations are represented by dashed circles. Neighboring particles have a black line drawn between them. . . . .	140
5.5	If $P$ is the topmost particle in a component of $m_i$ of size at least 2 and its neighborhood is connected, then (a)–(c) are the three possibilities for $N(P)$ . In all three of these cases, moving $P$ down-left satisfies Property 1. (d) and (e) show the two cases for subsequently moving $P$ to a new position such that the invariants still hold for $m_i$ . . . . .	141
5.6	The process of merging two lines stretching down-left from the same component of $m_{i+1}$ in order to satisfy Invariant 2. In (a) and (b), the moves must occur in the order listed. . . . .	143
5.7	100 particles in a line after (a) 1 million, (b) 2 million, (c) 3 million, (d) 4 million, and (e) 5 million iterations of $\mathcal{M}$ with bias $\lambda = 4$ . Edges have been drawn to indicate adjacencies between particles. . . . .	150
5.8	100 particles in a line after (a) 10 million and (b) 20 million iterations of $\mathcal{M}$ with bias $\lambda = 2$ . Edges have been drawn to indicate adjacencies between particles. . . . .	151
5.9	(a) The hexagonal lattice. (b) A self-avoiding walk in the hexagonal lattice. (c) A walk that is not self-avoiding. . . . .	153
5.10	(a) The duality between the triangular lattice and the hexagonal lattice. (b) An example of a particle configuration $\sigma$ , its corresponding polygon in the hexagonal lattice (shaded), and the self-avoiding walk bounding that polygon (bold). . . . .	153

5.11	All 11 connected hole-free configurations with three particles. In each the highest leftmost particle is labeled $H$ , and the lowest leftmost particle is labeled $L$ ; when there is only one leftmost particle $H = L$ . . . . .	157
5.12	The iterative process of Lemma 5.25. (a) One of the 11 connected hole-free configurations with three particle, and the two ways it can attach to the single particle with which the iterative process begins. (b) Another of the 11 connected hole-free configurations on three vertices, and the two ways it can attach to a configuration $\sigma$ with four particles. Particle adjacencies have been drawn as black lines, and dashed circles indicate unoccupied locations that guarantee no hole exists in the constructed configurations. . . . .	158
5.13	(a) In this image from [98], army ants of the genus <i>Eciton</i> build a dynamic bridge which balances the benefit of a shortcut path with the cost of committing ants to the structure. (b) Our shortcut bridging algorithm also balances competing objectives and converges to similar configurations. . . . .	171
5.14	Example instances $(L, O, \sigma_0, c, \alpha)$ of the shortcut bridging problem for which we present simulation results (Section 5.5). Light (brown) nodes are land locations, large (red) nodes are occupied by objects, and black nodes are occupied by particles. . . . .	172
5.15	A particle system using biases $\lambda = 4$ and $\gamma = 2$ to shortcut a V-shaped land mass with $\theta = \pi/3$ after (a) 2 million, (b) 4 million, (c) 6 million, and (d) 8 million iterations of Markov chain $\mathcal{M}$ , beginning in configuration $\sigma_0$ shown in Figure 5.14a. . . . .	182
5.16	A particle system using biases $\lambda = 4$ and $\gamma = 2$ to shortcut an N-shaped land mass after (a) 10 million and (b) 20 million iterations of Markov chain $\mathcal{M}$ , beginning in configuration $\sigma_0$ shown Figure 5.14b. . . . .	182
5.17	A particle system using biases $\lambda = 4$ and $\gamma = 2$ to shortcut a V-shaped land mass with angle (a) $\pi/6$ , (b) $\pi/3$ , and (c) $\pi/2$ after 20 million iterations of Markov chain $\mathcal{M}$ . For a given angle, the land mass $L$ and initial configuration $\sigma_0$ were constructed as described in Section 5.8. . . . .	183
5.18	The land mass $L$ of constant width 5 for (a) a small value of $\theta \sim \pi/6$ and height 8 and (b) a large value of $\theta \sim \pi/2$ and height 9. Point $m$ is the midpoint of the segment between the midpoints of $\ell_1$ and $\ell_2$ , and $b$ is shown as a dashed line. . . . .	184
5.19	The initial configuration $\sigma_0$ , with particles shown in black and objects enlarged and red, for (a) a small value of $\theta \sim \pi/6$ and (b) a large value of $\theta \sim \pi/2$ . Point $m$ is the midpoint of the segment between the midpoints of $\ell_1$ and $\ell_2$ , and $b$ is shown as a dashed line. . . . .	185

5.20 Figures from proofs in Section 5.8.1. (a) A depiction of the notation used in the proof of Lemma 5.44; the intersection of  $b_8$  and the gap is depicted as a solid segment, which is of length  $8\sqrt{3}\tan(\theta/2) + 1$  and contains 4 gap locations. (b) The configuration  $\sigma^*$  used in Lemma 5.45 for  $\theta = \pi/6$  and  $k = 8$ . . . . . 187

5.21 The path of length  $k$  (bold) from vertex  $v_1$  to the first land location in line  $b$  considered in the proof of Lemma 5.47; this path is used to calculate the gap height  $k$  in terms of the gap depth  $q$ . By also considering the reflection of this path from  $v_2$  (solid line), we can calculate the distance between the two objects to be  $q + 2\lceil w \rceil + 3$  (Lemma 5.48). . . . . 191

5.22 From the proof of Lemma 5.49: (a) An example of a shortest path between land locations on opposite sides of the gap passing through midpoint  $m$ . (b) The four possible locations for midpoint  $m$  for which a shortest path passing through or below  $m$  contains  $m'$ , and a shortest path from  $m'$  to a land location (solid line). . . . . 193

## SUMMARY

The problem of generating random samples from large, complex sets is widespread across the sciences, where such samples provide one way to begin to learn about the sets' typical properties. However, when the samples generated are unexpectedly correlated or drawn from the wrong distribution, this can produce misleading conclusions. One way to generate random samples is with *Markov chains*, which are widely used but often applied without careful analysis of their *mixing time*, how long they must run for until they are guaranteed to produce good samples. We present new mixing time bounds for two sampling problems from discrete geometry: *dyadic tilings*, combinatorial structures with applications in machine learning and harmonic analysis, and *3-colorings* on a grid, an instance of the celebrated antiferromagnetic Potts model from statistical physics. Both of these results required the development of new techniques.

In addition, we use Markov chains in a novel way to address research questions in programmable matter. Here, a main goal is to understand how simple computational elements can collectively accomplish complicated system-level goals. In an abstracted setting, we show that groups of particles executing our simple processes, based on Markov chains, can accomplish various tasks. This includes *compression*, a behavior exhibited by natural distributed systems such as fire ants and honey bees, and *shortcut bridging*, where the particles build bridges that optimize the same global trade-off as certain bridge-building ant colonies.

Throughout, a key ingredient is the interplay between global properties of Markov chains, including but not limited to mixing time, and their dependence on *local moves*, or Markov chain transitions that change only a small part of the configuration. We call the global behavior that arises out of these local moves and their probabilities *emergent behavior*. In addition to understanding the relationship between local moves and mixing times in order to give sampling guarantees, our work on programmable matter harnesses this interaction between local and emergent behavior in a novel way, to develop distributed algorithms.

# CHAPTER 1

## INTRODUCTION

The problem of generating a random sample from a large, complex set arises across many areas, such as polling [3], approximating statistics of real-world systems [63], and as a subroutine in randomized algorithms [32]. In these examples and more, studying random samples can tell us what a ‘typical’ element of a set looks like and provide insights about likely properties and behaviors. However, the problem of efficiently finding random elements is often a difficult one. One common approach uses a Markov chain: starting at an arbitrary configuration, iteratively make random local changes for long enough that, regardless of the starting point, we output a good random sample. This requires mathematically bounding the *mixing time*, the number of iterations until the configuration obtained is sufficiently random. Sampling algorithms using Markov chains are widespread throughout the natural and computational sciences, but are often applied without rigorous mixing time analysis, potentially producing misleading conclusions when samples are unexpectedly correlated or drawn from the wrong distribution. This thesis provides rigorous guarantees about the behavior of several Markov chain sampling algorithms, with a particular focus on problems from discrete geometry.

We also present a novel applications of Markov chains to programmable matter. Here our goal extends beyond generating random samples, and we develop decentralized, asynchronous algorithms for accomplishing various objectives in distributed systems. The intuition and analysis tools developed for understanding Markov chain sampling algorithms enabled this interdisciplinary effort to be successful and allows us to provide guarantees about the behavior of our distributed algorithms.

A common thread throughout this thesis is understanding the relationship between local and global behavior of Markov chains and their underlying models, which can be non-



intuitive and complex. Global properties of Markov chains are often affected by various (input) parameters that dictate the local behavior of the chains, and extremely small changes to some parameter – such as modifying the probabilities of certain moves – can have enormous impact. *Phase transitions* exist when there is a critical value for some parameter such that the behavior on either side of this critical point is radically different; for example, sampling algorithms may be prohibitively slow above a critical point but efficient otherwise. Furthermore, once we understand such emergent behavior, we can also harness it: understanding the relationship between probabilities of local moves and overall converge behavior of Markov chains is precisely the insight that enabled our work on programmable matter.

## 1.1 Random Sampling using Markov Chains

A main focus of this thesis is random sampling. How can you quickly choose an uniformly random element from a very large set? This is difficult when, for instance, even writing down everything in the set would take a prohibitively long amount of time.

Markov chains are one widely-used method for generating random samples. For example, the state-of-the-art algorithm for estimating the volume of high-dimensional convex bodies uses a Markov chain to take random samples from a sequence of high-dimensional distributions [32]. Markov chain sampling algorithms also appear as primitives within fast graph algorithms, such as in [52], and in machine learning algorithms, such as in [103]. More broadly, Markov chain Monte Carlo methods are used in a variety of statistical applications in the natural sciences. For example, such methods have been used in chemistry, to estimate degradation of chemicals in the soil [63]; in computational biology, to reconstruct phylogenetic trees [74]; and in linguistics, to infer probabilistic context-free grammars for spoken languages [72].

We begin with two examples of Markov chain sampling problems that are relevant to the work presented in this thesis and will illustrate some important points.

### 1.1.1 Card Shuffling: Global and Local Moves

A standard illustrative example of using a Markov chain to generate a random sample is shuffling a deck of cards. Here the set from which we wish to generate a random sample, the *state space*, consists of all possible orderings of the cards. The goal of shuffling is to quickly reach a random order on the deck of cards (a *random state* in this state space), by repeating random reordering/shuffling steps according to some iterative rules (making *random transitions* between states).

A common way cards are shuffled in practice is the *riffle shuffle*, where the deck of cards is divided into two parts and then the cards from each half are interleaved. This has been formalized, so that it can be studied mathematically, in the Gilbert-Shannon-Reeds model, which has been shown to be a good approximation of how humans perform the riffle shuffle (see [45], Chapter 4D). It is a well-known result that, for a certain notion of “close to uniform” that is often used when studying Markov chains (*total variation distance*), after seven riffle shuffles a deck of 52 is close to being in a uniformly random order. More generally, for decks of  $n$  cards, about  $1.5 \log_2 n$  riffle shuffles suffice to reach a nearly uniform random order [8]. Because a single riffle shuffle significantly changes the order of the cards in a deck, we call it a *global move*.

There are other methods for shuffling a deck of cards that just change one or two cards at a time; we call such transitions *local moves*. For many Markov chains, local moves are easier to identify and simpler to implement in practice. Examples of shuffling processes using local moves include the *top-to-random shuffle*, where the top card of the deck is placed at a random location in the deck, and the *random transposition shuffle*, where two randomly chosen cards are swapped. The number of times each of these shuffles has to be performed to guarantee a close to uniformly random ordering of the cards is  $\Theta(n \log_2 n)$  [2, 46]; it is impractical for a human to perform this many shuffling steps, but this process can easily be implemented and executed on a computer. Throughout this thesis a focus will be placed on Markov chains like these that use local moves: a single transition changes only a very

small part of the configuration. In many cases, we can design Markov chains using only local moves that as a whole still converge as desired, and often do so fairly quickly. The simplicity of such local processes is appealing. Furthermore, when transitions are local, often they can be implemented in a distributed way without requiring global knowledge of the whole system. This is a critical feature in the success of our stochastic, Markov chain-based approach to developing algorithms for distributed programmable matter systems (see Section 1.5 and Chapter 5).

### 1.1.2 The Hardcore Model from Statistical Physics

We now discuss the hardcore model from statistical physics as a second illustrative example of random sampling using a Markov chain. Related statistical physics models, the *Ising model* and *Potts model*, will both feature prominently in later chapters.

For a graph  $G$ , an *independent set* is a subset of the vertices of  $G$  such that no two vertices in the set are adjacent. One can make an analogy to physical systems where the vertices of  $G$  represent the possible locations for gas molecules, and the size of the molecules is large enough to preclude them from occupying adjacent locations. When graph  $G$  is a lattice, the independent set model is referred to as the hardcore lattice gas model; the term *hardcore* refers to the hard constraint that adjacent vertices cannot both be in the same independent set.

Computationally, we are interested understanding whether or not it is possible to efficiently sample a uniformly random independent set from a graph  $G$ . Answers to algorithmic questions about sampling from statistical physics systems are often closely tied to the properties of the systems themselves, such as how quickly correlations between spins at different sites decay with distance [13, 54]. One way to generate a random independent set is to use the following Markov chain: beginning at an arbitrary independent set, repeatedly pick a random vertex of  $G$ ; if the vertex is in the independent set, remove it; if the vertex is not in the independent set, add it if the resulting configuration is still independent. This

gives a simple way of moving between independent sets that only requires looking at some random local neighborhood of  $G$  in each step. These moves suffice to reach all independent sets, and in Chapter 2 we'll define the machinery necessary to prove that if you do enough of these random transitions, eventually you'll be equally likely to be at any independent set in  $G$ . Quantifying 'eventually,' the amount of time it takes to reach a random state, will be discussed more in Section 1.2.

## 1.2 Convergence Times of Markov Chains

To formalize the processes described above, a *Markov chain* is a process that iteratively makes random transitions between states in some state space and eventually outputs its current state as a random sample. For such a sampling process to be effective, we need to know that the sample output comes from a distribution that is close to our desired one, and that we reach such a distribution quickly. We give an introduction to these ideas here; more formal definitions can be found in Chapter 2.

Before we can say anything about the distribution from which we generate random samples, we must show that the random transitions we have defined suffice to reach all configurations in the state space. In many examples, such as for card shuffling and the hardcore model, this is nearly trivial, but for the Markov chains we consider in Chapter 5 it takes significant effort to show. If the transitions connect all states and additionally we show the state space is *aperiodic*, which is nearly always trivial to do, we say the Markov chain is *ergodic*. When a finite ergodic Markov chain process is executed for a large number of steps, it is known to eventually converge to a unique *stationary distribution* over the states, which is the distribution samples are being (approximately) generated from. That is, the probability of a particular state being output as the random sample is its probability in the stationary distribution. It is important to ensure that the stationary distribution of the Markov chain matches the distribution we wish to sample from, and tools (such as *detailed balance*) exist to verify this.

Bounding the time it takes for a Markov chain to reach its stationary distribution can be extremely challenging. We measure the distance from stationarity in terms of *total variation distance*, half the  $\ell_1$  distance between the stationary distribution and the probability distribution describing the state of the Markov chain at a given time. We say the chain has *mixed* when this total variation distance is less than  $\varepsilon$ ; as is standard, we often assume  $\varepsilon = 1/4$ . The *mixing time* is the amount of time it takes a Markov chain to become mixed from the worst-case starting state.

Often when a Markov chain is used in practice, it is run for long enough that samples “appear” to be random, but there are no theoretical guarantees on what distribution the sample is being drawn from and how close that distribution is to the stationary distribution. For this reason, rigorous bounds on the mixing time of Markov chains, and proof techniques to achieve them, are essential. There exist many well-established techniques for bounding mixing times – such as coupling [29], path coupling [16], comparison [43], decomposition [81, 83], and canonical paths [44, 107] – but there remain many problems for which these approaches fail to give meaningful mixing time bounds. Among other results, this thesis presents two new techniques for proving mixing time bounds, the bisection/block moves approach (Section 3.5) and randomized extensions (Chapter 4).

### 1.3 Biased Markov Chains

The examples presented in Sections 1.1.1 and 1.1.2 use Markov chains to generate samples drawn from a uniform distribution. In these and other settings, sometimes it can be more desirable to draw samples from biased distributions. Given a Markov chain with defined transitions and a target distribution over its state space, one can ensure the Markov chain converges to this target distribution by setting the probabilities of transitions according to a *Metropolis Filter* [88]. Changing the bias of a Markov chain (that is, modifying the probabilities of its transitions) can have profound impacts on the mixing behavior it exhibits, so this has to be carefully understood.

### 1.3.1 Biased Card Shuffling

Consider the following shuffling algorithm for a deck of  $n$  cards that are labeled  $1, 2, 3, \dots, n$ : choose two random adjacent cards in the deck and swap them with some probability. Changing the probability with which such a swap occurs changes the distribution over orderings of the deck that this shuffling process converges to as well as the time it takes to converge. When each such swap is made with probability  $1/2$ , this process still converges to the uniform distribution and does so in time  $\Theta(n^3 \log n)$  [118]. Suppose instead probabilities of adjacent swaps are set so that whenever cards  $i$  and  $j$  are picked with  $i < j$ , card  $i$  is placed before card  $j$  with probability  $p$  and after card  $j$  with probability  $1 - p$  for some fixed  $p > 1/2$ . It turns out that this process converges to a distribution that favors having the cards in order, and does so in time  $\Theta(n^2)$ , much faster than in the unbiased case [9]. There are also ways to set transition probabilities that cause the chain to converge slower. Bhakta et al. [10] give one such example, obtained via a reduction to biased lattice paths, where transition probabilities still favor putting the cards in order (it's more likely to put  $i$  before  $j$  when  $i < j$ ), but the shuffling process takes exponential time to converge.

### 1.3.2 The Weighted Hardcore Model from Statistical Physics

A similar dependence of convergence time on the probabilities of local moves can be found when considering a weighted version of the hardcore model presented in Section 1.1.2. Suppose, instead of sampling a uniformly random independent set of a graph  $G$ , we want to be more likely to sample the larger (more interesting) independent sets. We can add a bias  $\lambda > 1$  to the sampling process described above: pick a random vertex of  $G$ ; if the vertex is in the independent set, remove it with probability  $1/\lambda$ ; if the vertex is not in the independent set, add it if possible. As we are less likely to remove vertices from the independent set, it's reasonable to expect that our random sample is more likely to have more vertices, and this turns out to be the case.

It has been shown that the convergence time of this process depends critically on the value of  $\lambda$ . Even just for graphs  $G$  that are rectangular subsets of  $\mathbb{Z}^2$ , when  $\lambda$  is small it converges to its stationary distribution quickly [113], but when  $\lambda$  is large this takes exponential time [12]. It is conjectured that there is a critical point  $\lambda_c$  such that when  $\lambda$  is below this point the mixing time is fast and above it the mixing time is slow [7]. This is one example of a more general phenomenon known as a *phase transition*. Though this phase transition for the hard core model has been studied empirically, efforts to extend proofs of mixing behavior past the known bounds have been unsuccessful; currently, mixing is only known to be fast for  $\lambda < 2.48$  and slow for  $\lambda > 7.12$ , even though it is conjectured that  $\lambda_c \approx 3.79$ .

#### 1.4 Phase Transitions and Emergent Behavior

In this dissertation, a main focus is placed on global characteristics of Markov chains, including but not limited to mixing time, and how they depend on *local moves*, transitions between states that only change a small part of the configuration at a time. We refer to the global behavior that arises out of these local moves and their probabilities as *emergent behavior*. While understanding the relationship between local moves and mixing times is essential to ensuring widely-used sampling processes are reliable, this thesis also shows how to harness the relationship between local moves and emergent behavior in a novel way, to develop distributed algorithms for programmable matter.

In Chapter 3 we prove the existence of a phase transition for a natural local Markov chain on *dyadic tilings*: the chain has a parameter  $\lambda$  that determines the probabilities of local moves, and there is a critical point  $\lambda_c = 1$  such that for all  $\lambda < \lambda_c$  the convergence time is polynomial and for all  $\lambda > \lambda_c$  the convergence time is exponential. Furthermore, at the critical point  $\lambda = \lambda_c$  the convergence time is polynomial but a larger polynomial than when  $\lambda < \lambda_c$ . This behavior, where there is some bias parameter  $\lambda$  (*inverse temperature*) and the mixing time of a natural Markov chain is as fast as possible above some critical temperature, a larger polynomial at the critical temperature, and exponential below the

critical temperature, is conjectured in many systems, especially in statistical physics. There are very few instances for which this behavior has been rigorously confirmed. Notable examples include the Ising model on complete graphs [77, 48], regular trees [47], and two-dimensional lattice regions [78], and the Potts model on the complete graph [33] and the two-dimensional lattice [60], all of which required significant effort to analyze.

Chapter 4 considers the effect boundary conditions have on sampling from grid 3-colorings and lozenge tilings. While we expect sampling with a natural local Markov chain to be fast regardless of boundary conditions (i.e., there is only one phase), previous proofs were insufficient to show this. We were able to use a new technique to push past previous barriers and give efficient sampling results for even more boundary conditions. Of particular interest, the new regions we're now able to efficiently sample from include those needed for self-reducibility; for the first time, we can also approximately count 3-colorings of grid regions. These results move us significantly closer to rigorously verifying that for these problems the emergent behavior has no dependence on boundary conditions.

Our Markov chain algorithms for particle processes presented in Chapter 5 provably exhibit phased behavior that depends on the input to the problem, but these results are of a different flavor. For these problems, we wish to obtain a sample that with high probability has a certain property, and show how to bias a Markov chain to guarantee this is true at stationarity. Rather than different phases characterizing when a Markov chain enables efficient sampling and when it does not, in these algorithms different phases correspond to different interesting, desirable properties that the stationary distribution exhibits. Just by changing a single parameter, we can use the same algorithm to obtain radically different behaviors. While we conjecture a sharp phase transition between the two phases we see (as there is for dyadic tilings), current proof techniques are not sufficient to obtain such a result. The most significant contribution of this result is the novel connections we make between Markov chains for sampling from particle processes and distributed algorithms for programmable matter. The next section defines and explains programmable matter and



motivates our work on the topic.

## 1.5 Programmable Matter

To develop a system of *programmable matter*, one endeavors to create a material or substance that utilizes user input or stimuli from its environment to change its physical properties in a programmable fashion. Many such systems have been realized; a non-exhaustive list includes:

- *DNA computing*, where strands of DNA programmed with specific base sequences combine in solution to form specific arrangements [1];
- *Smart materials*, including 3D-printed wood that bends in a preprogrammed way when wet [34];
- *Modular robots* that can reconfigure themselves to accomplish different tasks, such as the ReBiS robot which can switch between bipedal and snake-like movement [109].
- *Swarm robotics*, where large groups of robots collectively perform tasks, like the kilobots of [102].

Programmable matter can be divided into *active* and *passive* types. In passive programmable matter systems, which includes most instances of DNA computing and smart materials, individual elements have little to no control over how they respond to their environment. Instead, they rely on physical properties and interactions with environmental conditions to produce the desired results. Much work in this area has focused on shape formation (e.g., [34, 108]), including a large body of research in molecular self-assembly (e.g., [30, 50, 116, 119]).

In contrast, in active programmable matter systems, individual computational units are capable of making decisions and acting on those decisions. For example, in self-reconfigurable modular robots each robotic module can adjust its connections to other

modules in order to form different structures [93], and in distributed swarms each robot makes independent decisions about what to do [102].

### 1.5.1 Self-organizing Particle Systems

We will focus on active programmable matter. Because instances of active programmable matter are incredibly varied, instead of focusing on just one system we will instead examine an abstraction that captures features that are common across many different active programmable matter systems. This will allow rigorous exploration of the general algorithmic capabilities and limitations of active programmable matter, free from the particular nuances and constraints of certain instantiations.

In a *self-organizing particle system*, individual units called *particles* with limited computational and communication abilities occupy the vertices and move along the edges of some graph (representing real space) in a distributed, asynchronous way [39]. We are interested in what these simple particles can collectively accomplish without any centralized control.

Initial work on self-organizing particle systems gave deterministic distributed algorithms for problems such as shape formation [42] and object coating [41]. When a small amount of randomness was added, leader election was also shown to be possible [35]. These carefully-constructed distributed algorithms provably accomplished exactly the outlined objectives, but a lack of robustness and a need for persistent memory limited their practicality for real programmable matter systems.

### 1.5.2 The Stochastic Approach

In [21], we initiated the stochastic approach to developing distributed algorithms for programmable matter. By basing our distributed algorithm on a Markov chain, we are able to provably accomplish our objectives in a robust, nearly-oblivious way. The viability of this approach, which we present in Chapter 5, was further validated in [4] and [22]. At a high

level, we begin with a Markov chain on particle arrangements that converges to a desired distribution using local moves. Because each local move changes only a very small part of the configuration – one particle moves to an adjacent location – and is independent of what is happening elsewhere, it is possible to implement these moves in a distributed way. Specifically, we can give a distributed algorithm for each particle to execute independently and asynchronously so that the desired collective behavior is still realized despite the lack of central control. Because our distributed algorithm comes from a Markov chain, we can leverage some of the many tools commonly used for Markov chain analysis to provide guarantees about the performance of the distributed algorithm. This novel application of Markov chains to programmable matter and distributed algorithms was made possible by intuition developed from studying Markov chain sampling algorithms from a rigorous perspective and a thorough understanding of how probabilities of local Markov chain moves determine its behavior.

### 1.5.3 Inspiration: Biological Systems

Many programmable matter systems are inspired by distributed natural systems that are able to do amazing things. Species of social insects offer perhaps the most immediate analogy to active programmable matter systems: individual ants, wasps, or bees are fairly simple organisms, but colonies as a whole exhibit remarkable collective behaviors that appear to be greater than the sum of their parts. For example, individual fire ants *Solenopsis invicta* struggle to swim in water, but when their nests are flooded they gather together to form rafts out of their own bodies and float for weeks to months until they find a new home [89]. When foraging, army ants of the genus *Eciton* build bridges that shortcut the distance other ants have to travel. The construction of these bridges seems to optimize a global trade-off between having shorter foraging paths but having fewer foraging ants because many are occupied in bridge-building. This is done despite each ant only having knowledge of its local environment, not the whole foraging trail [98]. Remarkable collec-

tive behaviors have also been observed in other species: honey bees choose hive locations based on decentralized recruitment [19] and cockroach larvae perform self-organizing aggregation using pheromones with limited range [68].

The programmable matter problems we solve in Chapter 5 take inspiration from simple behaviors observed in biological systems: compression, expansion, and bridge-building. Biologists don't entirely understand how such tasks are accomplished in nature, but our algorithms give one plausible explanation of how simple individual elements, such as ants, can collectively exhibit greater intelligence and accomplish difficult tasks.

## CHAPTER 2

### BACKGROUND

Formalizing many of the ideas discussed in Chapter 1 is necessary to discuss our results. A *Markov chain* is a memoryless random process on some state space  $\Omega$ ; in this thesis, we only consider Markov chains on finite discrete state spaces. In particular, a Markov chain randomly transitions between states of  $\Omega$  in a time-independent, or *stochastic*, fashion: there are fixed rules that prescribe the probabilities with which the chain transitions to its next state that depend only on its current state. The probabilities of moves have no dependence on any past behavior of the Markov chain, how long the chain has been running for, or anything else. We focus on discrete time Markov chains, where in each iteration of the Markov chain one transition occurs, though most of our results apply to the analogous continuous time chains as well. For a more detailed background on Markov chains than is included here, see [76].

Because of its stochasticity, we can completely describe a Markov chain by its transition matrix  $P$ , which is an  $|\Omega| \times |\Omega|$  matrix, indexed by the states of  $\Omega$ , defined such that for any pair  $x, y \in \Omega$ ,  $P(x, y)$  is the probability, if in state  $x$ , of moving to state  $y$  in one iteration of the Markov chain. The  $t$ -step transition probability  $P^t(x, y)$  is the probability of moving from  $x$  to  $y$  in exactly  $t$  steps.

A Markov chain is *irreducible* if there is a sequence of valid transitions from any state to any other state, that is, if for all  $x, y \in \Omega$  there is a  $t$  such that  $P^t(x, y) > 0$ . A Markov chain is *aperiodic* if for all  $x, y \in \Omega$ ,  $\gcd\{t : P^t(x, y) > 0\} = 1$ . A Markov chain is *ergodic* if it is both irreducible and aperiodic.

Any finite, ergodic Markov chain converges to a unique *stationary distribution*  $\pi$ , given by, for any  $x, y \in \Omega$ ,  $\pi(y) = \lim_{t \rightarrow \infty} P^t(x, y)$ ; importantly, this stationary distribution is completely independent of the starting state  $x$ . We can find the stationary distribution

of any finite ergodic Markov chain using *detailed balance*: any distribution  $\pi'$  that satisfies  $\pi'(x)P(x, y) = \pi'(y)P(y, x)$  for all  $x, y \in \Omega$  (the *detailed balance condition*) and has  $\sum_{x \in \Omega} \pi'(x) = 1$  must be the unique stationary distribution of the Markov chain (see, e.g., [55]). If a Markov chain satisfies the detailed balance equation, it is said to be *reversible*.

Given the transitions of a Markov chain and a desired stationary distribution  $\pi$  on  $\Omega$ , the Metropolis-Hastings algorithm [88] defines appropriate transition probabilities for the chain so that  $\pi$  is its stationary distribution. For a state  $x \in \Omega$ , we say its *neighbors* are the states it can transition to, and its *degree* is its number of neighbors. Starting at  $x \in \Omega$ , the Metropolis-Hastings algorithm picks a neighbor  $y \in \Omega$  uniformly with probability  $1/(2\Delta)$ , where  $\Delta$  is the maximum degree of any state, and moves to  $y$  with probability  $\min\{1, \pi(y)/\pi(x)\}$ ; with the remaining probability, it stays at  $x$  and repeats. If the given transitions connect the state space  $\Omega$  (i.e., if the chain is irreducible), then  $\pi$  must be the stationary distribution by detailed balance. While calculating  $\pi(y)/\pi(x)$  seems to require global knowledge, this ratio can often be calculated easily using only local information when many terms cancel out. This is the case for all biased local algorithms we consider in this thesis.

## 2.1 Convergence to Stationarity: Mixing Time, Spectral Gap and Relaxation Time

The time a Markov chain  $\mathcal{M}$  takes to converge to its stationary distribution  $\pi$  is typically measured in terms of the *total variation distance* between  $\pi$  and  $P^t(x, \cdot)$ , the probability distribution describing the state of  $\mathcal{M}$  at time  $t$  from starting state  $x$ ; formally,

$$\|P^t(x, \cdot), \pi\|_{TV} = \frac{1}{2} \sum_{y \in \Omega} |P^t(x, y) - \pi(y)|.$$

The *mixing time*  $t_{mix}$  of a Markov chain  $\mathcal{M}$  is the time it takes this total variation distance between  $P^t$  and the stationary distribution  $\pi$  to drop (and stay) below  $\varepsilon$ , from the worst

case starting state:

$$t_{mix}(\varepsilon) = \min \left\{ t : \forall t' \geq t, \max_{x \in \Omega} \left\| P^{t'}(x, \cdot), \pi \right\|_{TV} \leq \varepsilon \right\}.$$

As is standard, we often assume  $\varepsilon = 1/4$  and consider mixing time  $t_{mix} = t_{mix}(1/4)$ . We say  $\mathcal{M}$  is *rapidly mixing* if  $t_{mix}$  is bounded above by a polynomial in  $n$  and *slowly mixing* if it is bounded below by an exponential in  $n$ , where  $n$  is the size of the problem.

Closely related to mixing times of Markov chains are the notions of relaxation time and spectral gap. Following the notation of [76], the *absolute spectral gap*  $\gamma_*$  of a Markov chain  $\mathcal{M}$  with transition matrix  $P$  is  $1 - |\lambda_*|$ , where  $\lambda_*$  is the second largest eigenvalue of  $P$  in absolute value;  $P$ 's largest eigenvalue is always 1. The *relaxation time* of a reversible Markov chain is  $1/\gamma_*$ .

A *lazy* Markov chain is one where  $P(x, x) \geq 1/2$  for all  $x \in \Omega$ . For a lazy Markov chain  $\mathcal{M}$ , all eigenvalues are nonnegative, and so  $\lambda_* = \lambda_2$ , the second largest eigenvalue of  $P$ . The *spectral gap*  $\gamma$  is defined to be  $1 - \lambda_2$ , and for lazy Markov chains  $t_{rel} = 1/\gamma$ . In this thesis, we only examine the spectral gap and relaxation time for Markov chains that are lazy. The following well-known proposition relates the relaxation time and mixing time of a Markov chain; for a proof, see, e.g., [76, Theorem 12.3 and Theorem 12.4].

**Proposition 2.1.** *Let  $\mathcal{M}$  be an ergodic Markov chain on state space  $\Omega$  with reversible transition matrix  $P$  and stationary distribution  $\pi$ . Let  $\pi_{min} = \min_{x \in \Omega} \pi(x)$ . Then:*

$$(t_{rel} - 1) \log \left( \frac{1}{2\varepsilon} \right) \leq t_{mix}(\varepsilon) \leq \log \left( \frac{1}{\varepsilon \pi_{min}} \right) t_{rel}.$$

One can study the spectral gap of a Markov chain, and thus its relaxation and mixing times, by considering functions on the chain's state space. For  $f : \Omega \rightarrow \mathbb{R}$ , the *variance*

of  $f$  with respect to a distribution  $\pi$  on  $\Omega$  can be expressed as:

$$\text{var}_\pi(f) = \sum_{x \in \Omega} \pi(x) (f(x) - \mathbb{E}_\pi[f(x)])^2 = \frac{1}{2} \sum_{x, y \in \Omega} \pi(x)\pi(y)(f(x) - f(y))^2.$$

For a given reversible transition matrix  $P$  on state space  $\Omega$  with stationary distribution  $\pi$ , the *Dirichlet form*, also known as the *local variance*, associated to the pair  $(P, \pi)$  is, for any function  $f : \Omega \rightarrow \mathbb{R}$ ,

$$\mathcal{E}(f) = \frac{1}{2} \sum_{x, y \in \Omega} [f(x) - f(y)]^2 \pi(x) P(x, y).$$

The following well-known result (see, e.g., [76, Lemma 13.12]) demonstrates the close relationship between the spectral gap, Dirichlet form, and variance of a function on  $\Omega$ .

**Proposition 2.2.** *Given a Markov chain with reversible transition matrix  $P$  and stationary distribution  $\pi$ , the spectral gap  $\gamma = 1 - \lambda_2$  of  $P$  satisfies*

$$\gamma = \min_{\substack{f: \Omega \rightarrow \mathbb{R} \\ \text{var}_\pi(f) \neq 0}} \frac{\mathcal{E}(f)}{\text{var}_\pi(f)}.$$

In Chapter 3, in particular, we will bound mixing and relaxation times by appealing to spectral gaps and Dirichlet forms.

## 2.2 Techniques for Proving Mixing and Relaxation Time Upper Bounds

We now briefly summarize some well-known techniques that we will use to prove mixing and relaxation time upper bounds, first focusing on three different variants of the coupling approach. A *coupling* of a Markov chain  $\mathcal{M}$  is a joint Markov process  $(\mathcal{A}, \mathcal{B})$  on  $\Omega \times \Omega$  such that each of the marginals  $\mathcal{A}$  and  $\mathcal{B}$  is a faithful copy of  $\mathcal{M}$  and, once the two coordinates coalesce, they move in unison: if  $A_t = B_t$ , then  $A_{t+1} = B_{t+1}$ . There is a close relationship between the coupling time, the amount of time it takes for the two marginal chains in a



coupling to coalesce, and mixing and relaxation times of the chain.

### 2.2.1 Coupling

The first coupling theorem we state is a result of Chen [29] (see also [76, Theorem 13.1]).

**Theorem 2.3** ([29]). *Let  $\mathcal{M}$  be a Markov chain on state space  $\Omega$ , and let  $\Phi$  be a metric on  $\Omega$ . Suppose there exists a coupling  $(\mathcal{A}, \mathcal{B})$  of  $\mathcal{M}$  and a constant  $\beta < 1$  such that for any pair of states  $A_t, B_t \in \Omega$ , which are updated to  $A_{t+1}$  and  $B_{t+1}$ , respectively, after one iteration of the coupling, that*

$$\mathbb{E}[\Phi(A_{t+1}, B_{t+1}) \mid A_t, B_t] \leq \beta\Phi(A_t, B_t).$$

*Then, the spectral gap of  $\mathcal{M}$  satisfies  $\gamma \geq 1 - \beta$ .*

This will be used in Section 3.3 to give an upper bound on the relaxation time of the biased edge-flip Markov chain on dyadic tilings (for certain biases), and in Section 3.5.3 to prove that a certain Markov chain on dyadic tilings that uses block moves has a constant spectral gap.

### 2.2.2 Path Coupling

*Path coupling* arguments are a convenient way of bounding the mixing time of a Markov chain by considering only a subset  $U$  of the joint state space  $\Omega \times \Omega$  of a coupling. To show that  $\mathcal{M}$  is rapidly mixing, we consider an appropriate metric  $\Phi$  on  $\Omega$  and prove that the two marginal chains, if in a joint configuration in subset  $U$ , get no farther away in expectation after one iteration. Through linearity of expectation, this implies that any pair of configurations get no farther apart in one iteration of the coupling. Though path coupling was first proposed by Bubley and Dyer [16], we instead state the slightly cleaner version of the path coupling theorem due to Dyer and Greenhill [53].

**Theorem 2.4** ([53]). *Let  $\Phi : \Omega \times \Omega \rightarrow \mathbb{Z}$  be a metric which takes values in  $[0, S]$ . Let  $\mathcal{M}$  be an ergodic Markov chain on  $\Omega$  and let  $(\mathcal{A}, \mathcal{B})$  be a coupling of  $\mathcal{M}$ , with  $\Phi_t := \Phi(A_t, B_t)$  and  $\Delta\Phi_t := \Phi_{t+1} - \Phi_t$ . Let  $U \subseteq \Omega \times \Omega$  be such that for all  $(A_t, B_t) \in \Omega \times \Omega$ , there exists a path  $A_t = Z_0, Z_1, \dots, Z_r = B_t$  such that  $(Z_i, Z_{i+1}) \in U$  for  $0 \leq i < r$  and*

$$\sum_{i=0}^{r-1} \Phi(Z_i, Z_{i+1}) = \Phi(A_t, B_t).$$

*Suppose that for all  $(A_t, B_t) \in U$ , the coupling satisfies*

$$\mathbb{E}[\Delta\Phi_t \mid A_t, B_t] \leq 0.$$

*Additionally, assume there exists  $\alpha > 0$  such that for all  $t$  with  $\Phi_t \neq 0$ ,*

$$\mathbb{P}(\Delta\Phi_t \neq 0) > \alpha.$$

*Then the mixing time of  $\mathcal{M}$  satisfies*

$$t_{\text{mix}}(\varepsilon) \leq \left\lceil \frac{eS^2}{\alpha} \right\rceil \lceil \log(\varepsilon^{-1}) \rceil.$$

We use this path coupling theorem in Chapter 4 to show a Markov chain for grid 3-colorings mixes in polynomial time.

### 2.2.3 Multiplicative Path Coupling

Another variant of path coupling allows distances between states in  $\Omega$  to get exponentially large; with the previously stated path coupling theorem, such metrics would never yield polynomial time mixing bounds. To enable such a result, we need a stronger statement about the expected decrease in distance: the two coupled processes must, in expectation, get closer by a multiplicative factor in each iteration. A version of this theorem was first

proved in [16], but the statement presented in [62] that we give here will be easier for us to work with.

**Theorem 2.5** ([62]). *Let  $\phi : \Omega \times \Omega \rightarrow \mathbb{R}^+ \cup \{0\}$  be a metric which takes on finitely many values in  $\{0\} \cup [1, S]$ . Let  $\mathcal{M}$  be a lazy ergodic Markov chain on  $\Omega$  and let  $(\mathcal{A}, \mathcal{B})$  be a coupling of  $\mathcal{M}$ , with  $\phi_t := \phi(A_t, B_t)$ . Let  $U \subseteq \Omega \times \Omega$  be such that for all  $(A_t, B_t) \in \Omega \times \Omega$ , there exists a path  $A_t = Z_0, Z_1, \dots, Z_r = B_t$  such that  $(Z_i, Z_{i+1}) \in U$  for  $0 \leq i < r$  and*

$$\sum_{i=0}^{r-1} \phi(Z_i, Z_{i+1}) = \phi(A_t, B_t).$$

*Suppose there exists  $\beta < 1$  such that, for all  $(A_t, B_t) \in U$ ,*

$$\mathbb{E}[\phi_{t+1} \mid A_t, B_t] \leq \beta \phi_t.$$

*Then the mixing time of  $\mathcal{M}$  satisfies*

$$t_{\text{mix}}(\varepsilon) \leq \frac{\ln(S\varepsilon^{-1})}{1 - \beta}.$$

This result has also been referred to as the *exponential metric path coupling theorem*. We use this theorem to prove mixing time upper bounds for a biased Markov chain on dyadic tilings in Section 3.3.

#### 2.2.4 Comparison

Another method for obtaining a mixing time bound on a Markov chain is to compare the chain to similar Markov chains whose mixing time is already known. The comparison method was first used by Diaconis and Saloff-Coste [43], and has since been applied to a wide range of problems. Though Diaconis and Saloff-Coste state their comparison results in terms of Dirichlet forms, we go one step farther and state comparison results in terms of spectral gaps (as in Sections 13.4 and 13.5 of [76]).

Let  $E \subseteq \Omega \times \Omega$  be  $E = \{(x, y) \mid P(x, y) > 0\}$ . An  $E$ -path from  $x$  to  $y$  is a sequence  $\Gamma = (e_1, e_2, \dots, e_m)$  of edges in  $E$  such that  $e_i = (v_i, v_{i+1})$  for  $v_1, v_2, \dots, v_{m+1} \in \Omega$ , where  $v_1 = x$  and  $v_{m+1} = y$ . The length of an  $E$ -path is the number of edges it contains and is denoted  $|\Gamma|$ .

**Theorem 2.6.** *Let  $P$  and  $\tilde{P}$  be reversible transition matrices with stationary distributions  $\pi$  and  $\tilde{\pi}$ , respectively, on the same state space. For  $E = \{(x, y) \mid P(x, y) > 0\}$ , for each  $x$  and  $y$  such that  $\tilde{P}(x, y) > 0$  pick an  $E$ -path  $\Gamma_{x,y}$  from  $x$  to  $y$ . Let the congestion ratio be*

$$B = \max_{e=(z,w) \in E} \left( \frac{1}{\pi(z)P(z,w)} \sum_{x,y: \Gamma_{x,y} \ni e} \tilde{\pi}(x)\tilde{P}(x,y)|\Gamma_{x,y}| \right).$$

*Then the spectral gaps  $\gamma$  and  $\tilde{\gamma}$  of  $P$  and  $\tilde{P}$ , respectively, satisfy*

$$\tilde{\gamma} \leq \left[ \max_{x \in \Omega} \frac{\pi(x)}{\tilde{\pi}(x)} \right] B\gamma.$$

In particular, if two Markov chains with the same stationary distribution differ in their probabilities of moves by only a constant factor, then their spectral gaps, relaxation times, and mixing times also differ by at most a constant factor; such a situation is discussed in Section 4.2.5. Furthermore, while we prove results in Chapter 4 about a Markov chain that uses non-local tower moves, an application of the comparison technique, similar to that of [61] and [97], extends these fast mixing results to a local Markov chain known as Glauber dynamics that uses single site updates.

### 2.2.5 Other Techniques

Other techniques commonly used to prove mixing time upper bounds include canonical paths [81, 83], which can be viewed as a special case of the comparison approach, and decomposition [81, 83]. However, there remain many problems for which these approaches fail to give meaningful mixing time bounds. Because of this, new approaches, techniques,

and insights are needed. We make progress towards this goal in Section 3.5, where we use a new technique, based on work in statistical physics, to give a mixing time upper bound for the unbiased edge-flip Markov chain on dyadic tilings.

## 2.3 Techniques for Proving Mixing and Relaxation Time Lower Bounds

We complement many of the mixing and relaxation time upper bounds presented in this thesis with lower bounds. We briefly outline the three main techniques used to give mixing and relaxation time lower bounds; for more details on all of these approaches, see Chapter 7 of [76].

### 2.3.1 Diameter

The diameter of the state space of a Markov chain can be used to give a first lower bound on its mixing time.

**Theorem 2.7.** *Let  $D$  be the diameter of the state space of an irreducible aperiodic Markov chain. Then for any  $\varepsilon < 1/2$ , the mixing time satisfies*

$$t_{\text{mix}}(\varepsilon) \geq \frac{D}{2}.$$

We give a diameter lower bound for the biased edge-flip Markov chain for dyadic tilings in Section 3.3 to show that our mixing time upper bound is within a factor of  $n/\log n$  of the true mixing time of the chain.

### 2.3.2 Isoperimetric Inequalities

One can also give a lower bound on the time a Markov chain takes to converge to its stationary distribution by demonstrating that the state space contains a bottleneck. The expected time it takes the chain to cross this bottleneck is a lower bound on the mixing

time. We use this approach to give an exponential mixing time lower bound on the biased edge-flip Markov chain for dyadic tilings – for certain biases – in Section 3.4.

The *conductance* of an ergodic Markov chain  $\mathcal{M}$  with stationary distribution  $\pi$  is

$$\Phi_{\mathcal{M}} = \min_{\substack{S \subseteq \Omega \\ \pi(S) \leq 1/2}} \frac{1}{\pi(S)} \sum_{s_1 \in S, s_2 \in \Omega \setminus S} \pi(s_1) P(s_1, s_2).$$

We can infer a lower bound on the mixing time using the following theorem that relates conductance and mixing time (see, e.g., [106]).

**Theorem 2.8.** *For any Markov chain with conductance  $\Phi$ ,  $\forall \epsilon > 0$  we have*

$$t_{\text{mix}}(\epsilon) \geq \left( \frac{1}{4\Phi} - \frac{1}{2} \right) \log \left( \frac{1}{2\epsilon} \right).$$

### 2.3.3 Distinguishing Statistics

*Distinguishing statistics* are a generalization of conductance. Let  $f : \Omega \rightarrow \mathbb{R}$  be some function (a *statistic*) on the state space of some Markov chain  $\mathcal{M}$ . The partitions of  $\Omega$  into a set  $S$  and its complement  $\Omega \setminus S$  considered in the definition of conductance can be thought of a statistic where  $f(\sigma) = 1$  if  $\sigma \in S$  and  $f(\sigma) = 0$  if  $\sigma \notin S$ , but we can also consider more general statistics. For any such  $f$ , it follows from Proposition 2.2 that for a Markov chain  $\mathcal{M}$  with transition matrix  $P$  and stationary distribution  $\pi$  that  $\mathcal{M}$ 's spectral gap satisfies

$$\gamma \leq \frac{\mathcal{E}(f)}{\text{var}_{\pi}(f)}.$$

In Section 3.5.5, we give a distinguishing statistic for the unbiased edge-flip Markov chain that gives a spectral gap upper bound, and consequently a mixing and relaxation time lower bound. This result nicely complements our other results for the edge-flip Markov chain on dyadic tilings, and supports some interesting connections to statistical physics that we will discuss in Section 3.6.

## CHAPTER 3

### DYADIC TILINGS

In this chapter we examine the edge-flip Markov chain for dyadic tilings. A *dyadic tiling of size  $n$*  is a tiling of the unit square by  $n$  non-overlapping dyadic rectangles with the same area  $1/n$ , where a dyadic rectangle is one that can be written in the form  $[a2^{-s}, (a+1)2^{-s}] \times [b2^{-t}, (b+1)2^{-t}]$  for  $a, b, s, t \in \mathbb{Z}_{\geq 0}$ ; see Figure 3.1. More naturally, Lagarias, Spencer, and Vinson [73] showed that dyadic tilings are precisely those tilings that can be constructed by bisecting the unit square, either horizontally or vertically; bisecting each half again, either horizontally or vertically; and repeatedly bisecting all remaining rectangular regions until there are  $n$  total dyadic rectangles, each of equal area. We necessarily assume  $n$  is a power of 2.

Dyadic tilings have been used as a classifier in machine learning algorithms [104]; in harmonic analysis to approximate isotopic curves [18]; and their interesting combinatorial properties have been studied [73, 67]. More broadly, partitions of lattice regions into rectangles whose corners lie on lattice points, called *rectangular dissections*, arise in the study of VLSI layout [37], mapping graphs for floor layouts [94, 115], and routings and placements [120] and have long been of interest to combinatorialists [15, 112]. Of particular relevance to us are *equitable rectangular dissections* which require that all rectangles in the partition have the same area [64].

We are interested in understanding the efficiency of a natural method for generating random dyadic tilings via local moves, called the *edge flip* Markov chain. Given any dyadic tiling, this chain evolves by selecting an edge of the tiling uniformly at random and replacing it by its perpendicular bisector, if doing so yields a valid dyadic tiling of size  $n$ . A valid edge flip is shown in Figure 3.2a, while two edges that cannot be flipped are shown in Figures 3.2b and 3.2c. Similar edge flip Markov chains have been considered for domino




Figure 3.1: (a) A dyadic tiling of size 16 with a vertical bisector. (b) A dyadic tiling of size 16 with both a vertical and horizontal bisector. (c) A tiling that is not dyadic; the shaded rectangles are not dyadic because of their vertical dimension.




Figure 3.2: (a) A valid edge flip from one dyadic tiling to another. (b) Flipping the bold edge is not valid as the resulting tiling is not dyadic. (c) Flipping the bold edge is not valid as it does not result in a rectangular tiling.

tilings, where they have been shown to be rapidly mixing [79, 97, 118], and triangulations of point sets, where the mixing time of edge flips is a major open problem in computational geometry (see, e.g., [117]), though some special cases have been solved [87, 92].

The mixing time of the edge-flip Markov chain for dyadic tilings, and whether it mixes in polynomial time, was left as an open question by Janson, Randall, and Spencer in 2002 [67]. We answered this in the affirmative in [20], and we present these results in Section 3.5: we show the relaxation time of the edge flip Markov chain on dyadic tilings of size  $n$  is  $O(n^{4.09})$  and its mixing time is  $O(n^{5.09})$ .

As has been done for triangulations [27, 26], we also consider a weighted version of the edge-flip Markov chain for dyadic tilings. Here we bias our Markov chain so that it converges to a stationary distribution  $\pi$  where the probability of a tiling  $\sigma$  depends on the total length  $|\sigma|$  of all its edges:  $\pi(\sigma) \sim \lambda^{n|\sigma|}$ . When  $\lambda > 1$ , this favors dyadic tilings with



lots of long, thin rectangles. When  $\lambda < 1$ , this favors dyadic tilings with lots of squares and rectangles that are close to square. When  $\lambda = 1$ , this is exactly the unbiased edge-flip chain discussed above.

In [24], we showed that this weighted edge-flip Markov chain for dyadic tilings exhibits a phase transition at a critical point  $\lambda_c = 1$ . Specifically, we show that whenever  $\lambda < 1$ , the biased edge-flip Markov chain converges in polynomial time (Section 3.3) and whenever  $\lambda > 1$  it converges in exponential time (Section 3.4). Though phase transitions like this, where the behavior of a system changes dramatically at a single point, are conjectured to exist for many systems, proving they exist is rare.

More broadly, it is a general principle in statistical physics that in systems with some bias parameter (*temperature*) that induces different phases, the mixing time of natural heat-bath dynamics should be as fast as possible (the diameter of the state space) at high temperature, a larger polynomial at the critical temperature, and exponential at low temperature. However, there are very few instances for which this behavior has been rigorously confirmed. Exceptions are the Ising model on complete graphs [77, 48], regular trees [47], and the two-dimensional lattice [78], and the Potts model on the complete graph [33] and the two-dimensional lattice [60], all of which required significant effort to analyze. The edge-flip Markov chain for dyadic tilings is an example of heat-bath dynamics, and the parameter  $\lambda$  can be viewed as a function of inverse temperature. Our results confirm exponential mixing at low temperature ( $\lambda > 1$ ), polynomial mixing at high temperature ( $\lambda < 1$ ), and that the mixing time at the critical point ( $\lambda = 1$ ) is polynomial but strictly larger than the diameter of the state space (which is  $n \log(n)/2$ ), providing further evidence for this general statistical physics principle.

### 3.1 Preliminaries and Previous Work

A *dyadic interval* is an interval that can be written in the form  $[a2^{-s}, (a+1)2^{-s}]$  for non-negative integers  $a$  and  $s$  with  $0 \leq a < 2^s$ . As defined above, a *dyadic rectangle* is the

product of two dyadic intervals. A *dyadic tiling of size*  $n = 2^k$  is a tiling of the unit square by  $n$  dyadic rectangles of equal area  $1/n = 2^{-k}$  that do not overlap except on their boundaries; see Figure 3.1. Let  $\Omega_k$  be the set of all dyadic tilings of size  $n = 2^k$ . While we index by  $k$ , because there are no dyadic tilings unless  $n = 2^k$  for an integer  $k$ , our goal is still to give meaningful mixing time bounds in terms of  $n$ , the number of tiles.

We say a dyadic tiling has a *vertical bisector* if the line  $x = 1/2$  does not intersect the interior of any dyadic rectangle in the tiling. We say it has a *horizontal bisector* if the same is true of the line  $y = 1/2$ . The dyadic tiling in Figure 3.1a has a vertical bisector but no horizontal bisector, while the tiling in Figure 3.1b has both a vertical and horizontal bisector. It is easy to prove that every dyadic tiling of size  $n > 1$  has a horizontal bisector or a vertical bisector. We will also need the following lemma about dyadic intervals.

**Lemma 3.1.** *Two dyadic intervals of the same length do not overlap nontrivially.*

*Proof.* For any integer  $u$ , the only dyadic intervals of length  $2^{-u}$  are

$$[0, 2^{-u}], [2^{-u}, 2 \cdot 2^{-u}], [2 \cdot 2^{-u}, 3 \cdot 2^{-u}], \dots [k \cdot 2^{-u}, (k+1) \cdot 2^{-u}], \dots$$

and none of these overlap except at their endpoints. □

### 3.1.1 Combinatorics of Dyadic Tilings

The asymptotics of dyadic tilings were first explored by Lagarias, Spencer, and Vinson [73], and we present a summary of their results. Let  $A_k = |\Omega_k|$  denote the number of dyadic tilings of size  $n = 2^k$ . The unit square is the unique dyadic tiling consisting of one dyadic rectangle, so  $A_0 = 1$ . There are two dyadic tilings of size 2, since the unit square may be divided by either a horizontal or vertical bisector, so  $A_1 = 2$ . One can also observe that  $A_2 = 7$ ,  $A_3 = 82$ ,  $A_4 = 11047$ , ... . In fact, the values  $A_k$  can be shown to satisfy the recurrence  $A_k = 2A_{k-1}^2 - A_{k-2}^4$ ; we include a proof of this fact as presented in [67], because we will use these ideas later.

**Proposition 3.2** ([73]). *For  $k \geq 2$ , the number of dyadic tilings of size  $2^k$  is*

$$A_k = 2A_{k-1}^2 - A_{k-2}^4.$$

*Proof.* A dyadic tiling of size  $2^k$  has a horizontal bisector, a vertical bisector, or both. If it has a vertical bisector, the number of ways to tile the left half of the unit square is  $A_{k-1}$ ; by mapping  $x \rightarrow 2x$ , we can see that the left half of a dyadic tiling of size  $2^k$  is equivalent to a dyadic tiling of the unit square of size  $2^{k-1}$  because dyadic rectangles scaled by factors of two remain dyadic. Similarly, mapping  $x \rightarrow 2x - 1$ , the right half of a dyadic tiling of size  $2^k$  is equivalent to a dyadic tiling of size  $2^{k-1}$ . We conclude the number of dyadic tilings of size  $2^k$  with a vertical bisector is  $A_{k-1}^2$ . Similarly, by appealing to the maps  $y \rightarrow 2y$  and  $y \rightarrow 2y - 1$ , the number of dyadic tilings of size  $2^k$  with a vertical bisector is  $A_{k-1}^2$ . The number of dyadic tilings of size  $2^k$  with both a horizontal and a vertical bisector is  $A_{k-2}^4$ , as each quadrant of any such tiling is equivalent to a dyadic tiling of size  $2^{k-2}$ . This follows from appealing to the map  $(x, y) \rightarrow (2x, 2y)$  for the lower left quadrant, and appropriate translations of this for the other three quadrants. Altogether, we see  $A_k = A_{k-1}^2 + A_{k-1}^2 - A_{k-2}^4 = 2A_{k-1}^2 - A_{k-2}^4$ , as claimed.  $\square$

It is believed this recurrence does not have a closed form solution, but Lagarias, Spencer and Vinson proved the following about the number of dyadic tilings of a certain size.

**Lemma 3.3** ([73]). *The number of dyadic tilings of size  $n = 2^k$  satisfies  $A_k \sim \phi^{-1}\omega^{2^k}$ , where  $\phi = (1 + \sqrt{5})/2$  is the golden ratio and  $\omega = 1.84454757\dots$ ; an exact value for  $\omega$  is not known.*

We now define a recurrence for another useful statistic. We say that a dyadic tiling has a *left half-bisector* if the straight line segment from  $(0, 1/2)$  to  $(1/2, 1/2)$  doesn't intersect the interior of any dyadic rectangles. Figure 3.1a does not have a left half-bisector, while Figure 3.1b does. We are interested in the number of ways to tile the left half of a vertically-bisected dyadic tiling of size  $2^k$  such that it has a left half-bisector. Appealing to the dilation

maps defined in the proof of Proposition 3.2, this number is  $A_{k-2}^2$ . Among all possible ways to tile the left half of a vertically-bisected tiling  $\sigma \in \Omega_k$ , we define  $f_k$  to be the fraction with a left half-bisector. We see

$$f_k = \frac{A_{k-2}^2}{A_{k-1}}.$$

We can similarly define *right half-bisectors*, *top half-bisectors*, and *bottom half-bisectors* by considering the straight line segments between  $(1/2, 1/2)$  and, respectively,  $(1, 1/2)$ ,  $(1/2, 1)$ , and  $(1/2, 0)$ . Then  $f_k$  is also the fraction of tilings of the right half of vertically-bisected tiling  $\sigma$  with a right half-bisector, or the fraction of tilings of the top or bottom halves of a horizontally-bisected tiling  $\sigma$  with a top or bottom half-bisector, respectively. One can calculate  $f_2 = 0.5$ ,  $f_3 = 4/7 \sim 0.571$ , and  $f_4 = 49/82 \sim 0.598$ . We now examine the asymptotic behavior of  $f_k$ .

**Lemma 3.4.** *For all  $k \geq 3$ ,  $f_k = \frac{1}{2 - f_{k-1}^2}$ .*

*Proof.* This follows from the recurrence for  $A_k$  given in Proposition 3.2:

$$f_k = \frac{A_{k-2}^2}{A_{k-1}} = \frac{A_{k-2}^2}{2A_{k-2}^2 - A_{k-3}^4} = \frac{1}{2 - \frac{A_{k-3}^4}{A_{k-2}^2}} = \frac{1}{2 - f_{k-1}^2}.$$

□

We can use this recurrence to study the asymptotic behavior of the sequence  $\{f_k\}_{k=2}^\infty$ .

**Lemma 3.5.** *The sequence  $\{f_k\}_{k=2}^\infty$  is strictly increasing and bounded above by  $(\sqrt{5}-1)/2$ .*

*Furthermore,  $\lim_{k \rightarrow \infty} f_k = (\sqrt{5}-1)/2$ .*

*Proof.* Note  $f_2 = 0.5 < (\sqrt{5}-1)/2$ . Suppose by induction that  $f_{k-1} < \frac{\sqrt{5}-1}{2}$ . Then

$$f_k = \frac{1}{2 - f_{k-1}^2} < \frac{1}{2 - \left(\frac{\sqrt{5}-1}{2}\right)^2} = \frac{4}{8 - (6 - 2\sqrt{5})} = \frac{4}{2 + 2\sqrt{5}} = \frac{2}{1 + \sqrt{5}} = \frac{\sqrt{5}-1}{2}.$$

To show that  $f_k < f_{k+1}$  for all  $k \geq 2$ , it suffices to show  $x < 1/(2 - x^2)$  for all  $x \in [0.5, (\sqrt{5}-1)/2]$ . This is equivalent to showing the polynomial  $x^3 - 2x + 1$  is positive

in that range. Factoring shows this polynomial has roots at  $1$ ,  $(\sqrt{5}-1)/2$ , and  $-(\sqrt{5}+1)/2$ , and is positive in the range  $(-(\sqrt{5}+1)/2, (\sqrt{5}-1)/2)$ . This implies  $f_k < f_{k+1}$ , so the sequence is strictly increasing.

The sequence  $\{f_k\}_{k=2}^{\infty}$  is bounded and monotone, so it must converge to some limit  $\beta$ . To find  $\beta$ , we consider the function  $g(x) = 1/(2-x^2)$ , which is the recurrence for the  $f_k$ . This function is continuous away from  $\sqrt{2}$  and  $-\sqrt{2}$ , and thus certainly is continuous on  $[0.5, (\sqrt{5}-1)/2]$ , the range of possible values for the  $f_k$  and their limit  $\beta$ . This continuity implies

$$g(\beta) = g\left(\lim_{k \rightarrow \infty} f_k\right) = \lim_{k \rightarrow \infty} g(f_k) = \lim_{k \rightarrow \infty} f_{k+1} = \beta.$$

Thus the limit  $\beta$  is necessarily a fixed point of  $g(x)$ . The fixed points of  $g(x)$  are exactly the three roots of  $x^3 - 2x + 1$  found above, and the only one in  $[0.5, (\sqrt{5}-1)/2]$  is  $(\sqrt{5}-1)/2$ . We conclude  $\lim_{k \rightarrow \infty} f_k = (\sqrt{5}-1)/2$ , as desired.  $\square$

This recurrence for  $f_k$  will be used throughout various proofs in Section 3.5.

### 3.1.2 Markov Chains on Dyadic Tilings

While the previous subsection gave necessary preliminaries on dyadic tilings, our focus will be on Markov chains, not combinatorics. In 2002, Janson, Randall and Spencer were the first to consider Markov chains on dyadic tilings [67]. They proposed the unbiased edge-flip Markov chain, which uses local moves, and they showed it is irreducible but left as an open problem to derive that the mixing time is polynomial in  $n$ ; we solve this open problem in Section 3.5. Additionally, they presented a different nonlocal Markov chain on dyadic tilings, which has additional global moves consisting of rotations at all scales, and showed that this chain mixes in polynomial time. Specifically, their chain picks a random dyadic subrectangle of the unit square (of any area) and if no tiles cross the boundary of this region, it randomly rotates the tiling within in by 90, 180, or 270 degrees, rescaling as necessary. However, applications of the comparison technique of Diaconis and Saloff-

Coste [43] have failed to extend this polynomial mixing bound to the more natural local edge-flip Markov chain (which, in fact, corresponds to only performing rotations at the smallest scale).

In [24], we were the first to consider weighted Markov chains for dyadic tilings. Just as Caputo, Martinelli, Sinclair, and Stauffer did for triangulations [27], we weight dyadic tilings according to the total length of all the edges in the tiling. The authors of [27] conjectured a phase transition for triangulations, but were unable to give proofs that guaranteed its existence. We see this same phase transition in dyadic tilings, and are able to prove it occurs. We formally define this weighted Markov chain next; because of the locality of edge-flip moves, it is straightforward using a Metropolis Filter [88] to ensure the edge-flip Markov chain converges to the desired weighted distribution.

### 3.2 The Edge-flip Markov Chain on Weighted Dyadic Tilings

Let  $n = 2^k$ . For  $k \geq 1$ , the edge-flip Markov chain  $\mathcal{M}_k^{edge}$  on the state space  $\Omega_k$  of all dyadic tilings of the unit square of size  $2^k$  is given by Algorithm 1.

---

**Algorithm 1** Weighted Edge-Flip Markov Chain  $\mathcal{M}_k^{edge}$  on Dyadic Tilings of size  $n = 2^k$

---

Beginning at any  $\sigma_0 \in \Omega_k$ , repeat:

- 1: Choose a rectangle  $R$  of  $\sigma_i$  uniformly at random.
  - 2: Choose *left*, *right*, *top*, or *bottom* uniformly at random; let  $e$  be the corresponding side of  $R$ .
  - 3: Choose a uniformly random  $p \in (0, 1)$ .
  - 4: **if**  $e$  bisects a rectangle of area  $2^{-k+1}$ ; replacing  $e$  with its perpendicular bisector  $f$  yields a valid dyadic tiling; and  $p < \lambda^{n(|f|-|e|)}$  **then**
  - 5:      $\sigma_{i+1}$  is obtained from  $\sigma_i$  by replacing  $e$  with  $f$ .
  - 6: **else**  $\sigma_{i+1} = \sigma_i$ .
- 

We note this differs slightly from the way we originally presented this weighted Markov chain in [24]; in that paper, we considered dyadic tilings of size  $n$  that were scaled up to occupy an  $n \times n$  rectangle, rather than the unit square. When calculating edge distances in Step 4 of Algorithm 1, we correspondingly scale these edge distance up by a factor of  $n$  to match the probabilities used in the original weighted chain of [24].

We now consider the properties of  $\mathcal{M}_k^{edge}$ .

**Lemma 3.6** ([67]). *Markov chain  $\mathcal{M}_k^{edge}$  is irreducible.*

*Proof.* Connectivity of the state space  $\Omega_k$  follows from work on dyadic tilings in [67], specifically from their tree representation of a dyadic tiling. Dyadic constraints ensure rectangles exist in pairs, meaning there is always a possible edge-flip move for every rectangle. Using this fact, from any tiling, one can always reach the tiling  $\sigma_v$  consisting only of vertical  $1/n \times n$  rectangles by repeatedly finding edge-flips that turn a horizontal edge into a vertical edge. Because all moves are reversible ( $P(\sigma, \tau) > 0$  if and only if  $P(\tau, \sigma) > 0$ ), there also exists a sequence of valid moves turning  $\sigma_v$  into any other tiling. By going through  $\sigma_v$ , it is possible to use edge-flip moves to go between any two tilings in  $\Omega_k$ , meaning edge-flip moves connect  $\Omega_k$  and thus  $\mathcal{M}_k^{edge}$  is irreducible on this state space.  $\square$

**Lemma 3.7.** *Markov chain  $\mathcal{M}_k^{edge}$  is lazy.*

*Proof.* For any rectangle  $R$  of a dyadic tiling at most one of its left and right edges can be flipped to produce another valid dyadic tiling. This is because if  $R$ 's projection onto the  $x$ -axis is dyadic interval  $[a2^{-s}, (a+1)2^{-s}]$  for  $a, s \in \mathbb{Z}_{\geq 0}$ , then flipping its left edge yields a rectangle with  $x$ -projection  $[(a-1)2^{-s}, (a+1)2^{-s}]$  and flipping its right edge yields a rectangle with  $x$ -projection  $[a2^{-s}, (a+2)2^{-s}]$ . If  $a$  is even, the first of these intervals is not dyadic, while if  $a$  is odd, the second is not, so at most one of these edge flips produces a valid dyadic tiling. Similarly, at most one of  $R$ 's top and bottom edges yields a valid edge flip. This implies in each iteration with probability at least  $1/2$  a pair  $(R, e)$  is selected that does not yield a valid edge flip move.  $\square$

**Lemma 3.8.** *Markov chain  $\mathcal{M}_k^{edge}$  on dyadic tilings of size  $n = 2^k$  with parameter  $\lambda$  has a unique stationary distribution  $\pi$  given by, for  $\sigma \in \Omega_k$ ,*

$$\pi(\sigma) = \frac{\lambda^{n|\sigma|}}{Z},$$

where  $|\sigma|$  is the sum of the lengths of all the edges in tilings  $\sigma$  and  $Z = \sum_{\sigma \in \Omega_k} \lambda^{n|\sigma|}$  is the normalizing constant, also called the partition function.

*Proof.* By Lemmas 3.6 and 3.7,  $\mathcal{M}_k^{edge}$  is irreducible and aperiodic and thus is ergodic. Because  $\mathcal{M}_k^{edge}$  is ergodic on a finite state space, it converges to a unique stationary distribution. Distribution  $\pi$  given in the statement of the lemma satisfies detailed balance: for tilings  $\sigma$  and  $\tau$  that differ by a flip of an edge  $e$  in  $\sigma$  to an edge  $f$  in  $\tau$ , we see that Algorithm 1 yields

$$P(\sigma, \tau) = \frac{\min(1, \lambda^{n(|f|-|e|)})}{4n},$$

$$P(\tau, \sigma) = \frac{\min(1, \lambda^{n(|e|-|f|)})}{4n}.$$

If  $|f| \geq |e|$ , then  $P(\sigma, \tau) = \lambda^{n(|f|-|e|)}/4n$  and  $P(\tau, \sigma) = 1/4n$ , and similarly we see

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{n|\sigma|}}{Z} \cdot \frac{\lambda^{n(|f|-|e|)}}{4n} = \frac{\lambda^{n(|\sigma|-|e|+|f|)}}{Z \cdot 4n} = \frac{\lambda^{n|\tau|}}{Z} \frac{1}{4n} = \pi(\tau)P(\tau, \sigma).$$

If  $|e| \geq |f|$ , then  $P(\sigma, \tau) = 1/4n$  and  $P(\tau, \sigma) = \lambda^{n(|e|-|f|)}/4n$ . In this case detailed balance is also satisfied:

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{n|\sigma|}}{Z} \cdot \frac{1}{4n} = \frac{\lambda^{n(|\sigma|-|f|+|e|)}}{Z \cdot 4n} = \frac{\lambda^{n|\tau|}}{Z} \frac{\lambda^{n(|e|-|f|)}}{4n} = \pi(\tau)P(\tau, \sigma).$$

Thus  $\pi$  satisfies detailed balance. Finally, we check that  $\pi$  is indeed a probability distribution:

$$\sum_{\sigma \in \Omega_k} \pi(\sigma) = \sum_{\sigma \in \Omega_k} \frac{\lambda^{n|\sigma|}}{Z} = \sum_{\sigma \in \Omega_k} \frac{\lambda^{n|\sigma|}}{\sum_{\sigma \in \Omega_k} \lambda^{n|\sigma|}} = 1.$$

We conclude  $\pi$  is the unique stationary distribution of  $\mathcal{M}_k^{edge}$ , as claimed.  $\square$

We implemented Markov chain  $\mathcal{M}_k^{edge}$  and simulated its behavior for a variety of values of  $k = \log n$  and  $\lambda$ , beginning at the tiling  $\sigma_v$  consisting of all  $1/n \times 1$  vertical rectangles. The result of such simulations for  $k = 6$ ,  $n = 64$ , and  $\lambda = 0.8, 1$ , and  $1.03$  after 1,000,000 iterations are shown in Figure 3.3. Consistently, when  $\lambda < 1$  we see that  $\mathcal{M}_k^{edge}$  converges






Figure 3.3: State of Markov chain  $\mathcal{M}_k^{edge}$  with  $k = 6$  and  $n = 64$  and various values of  $\lambda$  after 1,000,000 simulated steps, starting from a tiling with all  $1/n \times 1$  vertical rectangles.

quickly to a distribution that favors tilings with square or nearly-square rectangles. When  $\lambda > 1$  even after many steps there remain tilings with many vertical rectangles, a clear lingering dependence on the initial state that indicates slow convergence. When  $\lambda = 1$  the chain also seems to converge quickly, but there is not such a strong preference for rectangles that are square or nearly square as there is when  $\lambda < 1$ . In the next sections we verify all of these observations rigorously and show that  $\mathcal{M}_k^{edge}$  exhibits a phase transition at critical point  $\lambda_c = 1$ .

### 3.3 Polynomial Convergence when $\lambda < 1$

When  $\lambda < 1$ , the stationary distribution of  $\mathcal{M}_k^{edge}$  favors tilings that have more rectangles that are square or close to square. We use this bias in a critical way in our proofs to show that  $\mathcal{M}_k^{edge}$  mixes in polynomial time –  $O(n^2)$  – for all  $\lambda < 1$ . The biased edge-flip Markov chain on grid triangulations discussed above similarly favors triangles with low aspect ratio at stationarity whenever  $\lambda < 1$  [27]. Simulations suggest that their chain mixes in polynomial time for these values of  $\lambda$ , but they are only able to prove fast mixing for  $\lambda < \lambda_0$  for a small constant  $1/8 \leq \lambda_0 < 1/4$ .<sup>1</sup> In contrast, our results apply for all  $\lambda < 1$ , which, with the complementary results of slow mixing for  $\lambda > 1$  in the next section, guarantees the existence of a phase transition at  $\lambda_c = 1$  only conjectured but not verified for triangulations.

---

<sup>1</sup>The authors of [27] do not give an exact value for  $\lambda_0$  but these bounds can be extracted from their proofs.

Specifically, we prove that for all even  $k$ ,  $\mathcal{M}_k^{edge}$  is rapidly mixing whenever  $\lambda < 3^{-1/\sqrt{n}}$  (Theorem 3.12). This bound approaches 1 as  $n$  grows, so for any  $\lambda < 1$  there is sufficiently large  $n$  for which the Markov chain  $\mathcal{M}_n$  is rapidly mixing (Theorem 3.13). To give some perspective, we note that for all  $n \geq 4$  ( $k \geq 2$ ), we have fast mixing for all  $\lambda < 0.577$ . Already for  $n \geq 1024$  ( $k \geq 10$ ) we have fast mixing for all  $\lambda < 0.966$ . That  $k$  is even implies  $n$  is a perfect square and there exists a “ground state” tiling consisting entirely of  $(1/\sqrt{n}) \times (1/\sqrt{n})$  squares, necessary for the proofs in this section.

We use a multiplicative path coupling argument with an exponential metric. Though multiplicative path coupling was first considered in [16], we use the version of the path coupling theorem as stated in [62], which will be easier to apply for our purposes. In a path coupling argument, instead of arguing about the coalescence of two coupled chains in an arbitrary pair of states, it suffices to assume the configurations of the two coupled chains are a pair of states in some subset  $U \subset \Omega_k \times \Omega_k$ . For our purposes,  $U$  consists of all pairs of tilings of size  $n = 2^k$  that differ by one edge-flip (i.e., that are adjacent in  $\Omega_k$ ). We will show that for any coupling whose joint state is two configurations in  $U$ , after one iteration of the Markov chain, the expected distance between the two coupled chains decreases by a constant factor of their original distance. It is crucial to define the appropriate notion of “distance” between two tilings; we do so by carefully defining the distance between tilings that are adjacent in  $\Omega_k$ , and then extending this definition to non-adjacent tilings.

**Definition 3.9.** *Consider any dyadic tilings  $\sigma_1$  and  $\sigma_2$  that differ by one flip between edge  $e$  and edge  $f$ , both bisecting a common area  $2/n$  rectangle  $S$ . Without loss of generality, suppose that  $|e| \geq |f|$ . We define the distance between  $\sigma_1$  and  $\sigma_2$  to be*

$$\phi(\sigma_1, \sigma_2) = \phi(\sigma_2, \sigma_1) := \lambda^{n(|f|-|e|)},$$

*For any dyadic tilings  $\sigma_1$  and  $\sigma_2$  that are not adjacent in  $\Omega_k$ , the distance between  $\sigma_1$  and  $\sigma_2$  is the minimum over paths in  $\Omega_k$  from  $\sigma_1$  to  $\sigma_2$  of the sum of the distances between*




Figure 3.4: A sequence of edge-flips from tiling  $\sigma_h$  consisting of all  $1 \times 1/n$  horizontal rectangles to ground state tiling  $\sigma^*$  consisting of all  $(1/\sqrt{n} \times 1/\sqrt{n})$  rectangles.

*adjacent tilings along the path.*

*If  $\sigma_1 = \sigma_2$ , then  $\phi(\sigma_1, \sigma_2) = 0$ .*

We now examine the range of this distance metric.

**Lemma 3.10.** *For  $\lambda < 1$ , the distance metric  $\phi$  of Definition 3.9 takes on values in the range  $0 \cup [1, n \log(n)\lambda^n]$ .*

*Proof.* If  $\sigma_1$  and  $\sigma_2$  differ by an edge-flip, then  $\phi(\sigma_1, \sigma_2)$  is at least 1 as  $\lambda < 1$  is being raised to a nonpositive power. If  $\sigma_1$  and  $\sigma_2$  differ by more than an edge-flip, then  $\phi(\sigma_1, \sigma_2)$  is a sum of distances that are at least 1 so is at least 1. It only remains to show that  $\phi(\sigma_1, \sigma_2)$  has the stated upper bound.

Let  $\sigma^*$  denote the ground state tiling, tiling the unit square with  $n$  smaller squares of size  $(1/\sqrt{n}) \times (1/\sqrt{n})$ . Careful consideration shows that the two dyadic tilings at farthest distance  $\phi$  from  $\sigma^*$  are the tiling consisting of all  $1 \times 1/n$  horizontal rectangles  $\sigma_h$  and the tiling consisting of all  $1/n \times 1$  vertical rectangles  $\sigma_v$ . We note that one path in  $\Omega_k$  from  $\sigma_h$  to  $\sigma^*$  consists of  $(\log n)/2 = k/2$  stages, where in each stage  $n/2$  edge-flips are performed, reducing the length of each of the  $n$  rectangles by half; see Figure 3.4.

The contribution to  $\phi(\sigma_h, \sigma^*)$  from each of these edge-flips is at most  $\lambda^{-n}$ , as for any two tilings differing by a flip of edge  $e$  to edge  $f$  we have that  $||e| - |f|| \leq 1$ , so raising  $\lambda^{-1} > 1$  to  $n$  times this amount is at most  $\lambda^{-n}$ . There are  $nk/4$  such moves in this particular path in  $\Omega_k$  from  $\sigma_h$  to  $\sigma^*$ , giving  $\phi(\sigma_h, \sigma^*) \leq (nk/4)\lambda^{-n}$ . The same holds for  $\sigma_v$ . There is

thus a path between any two tilings, through the ground state  $\sigma^*$ , yielding the bound

$$\phi(\sigma_1, \sigma_2) \leq (nk/2)\lambda^{-n} \leq n \log(n)\lambda^{-n}.$$

Thus  $\phi$  takes on values in the range  $\{0\} \cup [1, n \log(n)\lambda^{-n}]$ , as claimed.  $\square$

To formally define our coupling of Markov chain  $\mathcal{M}_k^{edge}$ , we first restate  $\mathcal{M}_k^{edge}$  in a slightly different way, making it explicitly lazy and easier to couple; see Algorithm 2.

---

**Algorithm 2** Weighted Edge-Flip Markov Chain  $\mathcal{M}_k^{edge}$  on Dyadic Tilings of size  $n = 2^k$

---

Beginning at any  $\sigma_0 \in \Omega_k$ , repeat:

1: Choose, uniformly at random,  $(x, y, d, o, p) \in$

$$\left\{ \frac{1}{2n}, \frac{3}{2n}, \frac{5}{2n}, \dots, \frac{2n-1}{2n} \right\} \times \left\{ \frac{1}{2n}, \frac{3}{2n}, \frac{5}{2n}, \dots, \frac{2n-1}{2n} \right\} \times \{t, l, b, r\} \times \{0, 1\} \times (0, 1).$$

2: Let  $R$  be the (unique) rectangle in  $\sigma_t$  containing  $(x, y)$ .

3: If  $d = t$ , let  $e$  be the top boundary of  $R$ ; if  $d = l, b$ , or  $r$ , let  $e$  be the left, bottom, or right boundary of  $R$ , respectively.

4: **if**  $e$  bisects a rectangle of area  $2^{-k+1}$ ;  $\log |e| \equiv o \pmod{2}$ ; replacing  $e$  with its perpendicular bisector  $f$  yields a valid dyadic tiling; and  $p < \lambda^{n(|f|-|e|)}$  **then**

5:  $\sigma_{i+1}$  is obtained from  $\sigma_i$  by replacing  $e$  with  $f$ .

6: **else**  $\sigma_{i+1} = \sigma_i$ .

---

We note that Algorithm 2 is simply Algorithm 1 with an additional stationary probability at each iteration. In particular, each rectangle  $R$  of any tiling  $\sigma$  is of area  $1/n$  and so contains exactly  $n$  points in the regular  $n^2$ -sized grid  $\left\{ \frac{1}{2n}, \frac{3}{2n}, \frac{5}{2n}, \dots, \frac{2n-1}{2n} \right\} \times \left\{ \frac{1}{2n}, \frac{3}{2n}, \frac{5}{2n}, \dots, \frac{2n-1}{2n} \right\}$ , meaning the probability that a given rectangle is chosen by Algorithm 2 is exactly  $1/n$ , the same probability with which each rectangle is chosen in Step 1 of Algorithm 1. A given flippable edge  $e$  in  $\sigma$  is thus selected by  $2n$  different values of  $(x, y, d, o)$ , specifically, the  $2n$  points  $(x, y)$  in the two rectangles  $e$  separates, each with the appropriate value of direction  $d$  and parity  $o$ . Consequently, a given flippable edge  $e$  is selected by  $(x, y, d, o)$  with probability  $2n \cdot \frac{1}{n^2} \cdot \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{4n} =: q$ . In Algorithm 1, each edge  $e$  is selected by 2 choices of  $R, e$ , a selection which occurs with probability  $2/(n \cdot 4) = 1/2n$ . In both version of  $\mathcal{M}_k^{edge}$ , this flip then occurs with probability  $\min\{1, \lambda^{n|\sigma'|-n|\sigma|} = \lambda^{n(|f|-|e|)}\}$ , according to




Figure 3.5: Rectangle  $S$  of area  $2/n$  in marginal tilings (a)  $A_t$  and (b)  $B_t$ .

the random value of  $p$ . This change to the stationary probability of  $\mathcal{M}_k^{edge}$  will have no effect on its asymptotic mixing time, but will make our coupling argument much easier; a coupling argument for Algorithm 1 is possible, but requires unnecessary technical details that can be avoided by instead working with Algorithm 2, as we do throughout this section.

Let  $(\mathcal{A}, \mathcal{B})$  denote a coupling of Markov chain  $\mathcal{M}_k^{edge}$  on dyadic tilings, as stated in Algorithm 2. Let  $A_t$  and  $B_t$  denote the states of these two coupled chains, respectively, after  $t$  iterations. At each iteration,  $A_t$  and  $B_t$  are simultaneously updated according to Algorithm 2 by choosing the same random values  $(x, y, d, o, p)$  for each. Let  $\phi_t = \phi(A_t, B_t)$  denote the distance between the two chains in the coupling  $(\mathcal{A}, \mathcal{B})$  after  $t$  iterations. We now show that if  $A_t$  and  $B_t$  differ by a single edge-flip, after one iteration of  $\mathcal{M}_k^{edge}$  in expectation they are closer together. By linearity of expectation, even if  $A_t$  and  $B_t$  differ by more than one edge-flip, after one iteration of  $\mathcal{M}_k^{edge}$  they are closer in expectation; this is the crux of the path coupling approach.

**Lemma 3.11.** *Suppose  $A_t$  and  $B_t$  differ by a single edge flip. Then for all  $\lambda < 3^{-1/\sqrt{n}}$ , there is a constant  $c$  such that for  $q = 1/(4n)$ ,*

$$\mathbb{E}[\phi_{t+1} \mid A_t, B_t] \leq (1 - qc)\phi_t$$

*Proof.* Let  $A_t$  and  $B_t$  differ by a single flip between edge  $e$  and edge  $f$ , where without loss of generality  $|e| \geq |f|$ ,  $e$  is horizontal in  $A_t$  of length  $2a$ ,  $f$  is vertical in  $B_t$  of length  $2b$ , and both bisect a rectangle  $S$  of area  $2/n$ ; see Figure 3.5. We wish to bound  $\mathbb{E}[\phi_{t+1} - \phi_t]$  in terms of  $\phi_t$ .

Any potential moves  $(x, y, d, o, p)$  that select an edge not in  $S$  and not on the boundary of  $S$  have the same effect on both  $A_t$  and  $B_t$  and thus, in these cases,  $\phi_{t+1} = \phi_t$ , as  $A_{t+1}$  and  $B_{t+1}$  still differ by the same single edge flip.

Because there is a rectangle in valid dyadic tiling  $A_t$  of dimension  $2a \times b$ , this implies that  $2ab = 1/n = 2^{-k}$ . As  $a$  and  $b$  are (negative) powers of 2,  $a \geq b$  by assumption, and  $k$  is even, then  $a = 2^i b$  where  $i$  is positive and odd. We now consider two cases,  $a \geq 8b$  and  $a = 2b$ .

**Case  $a \geq 8b$ .** We first examine the moves that decrease the distance between the two coupled chains. There are exactly two edge-flips that do this, namely flipping  $e$  to  $f$  in  $A_t$  or flipping  $f$  to  $e$  in  $B_t$ . There are  $2n$  values of  $(x, y, d, o)$  that select edge  $e$  in  $A_t$ . Precisely, these are each of the  $2n$  points  $(x, y)$  in  $S$  together with the appropriate direction from among  $t, b$  that selects  $e$  and the appropriate parity  $o$  such that  $\log |e| = o \pmod{2}$ . Examining the parity  $o$  shows these same choices do not yield a flippable edge in  $B_t$ ; this is where the value of  $o$  plays a critical role, as no edges within or on the top or bottom boundary of  $S$  in  $B_t$  have the same length as  $e$ . As each such selection of  $(x, y, d, o)$  occurs with probability  $1/(8n^2)$ , potential edge-flip  $e$  is selected with probability  $q = 1/(4n)$ . In this case the condition for flipping edge  $e$  is  $p < \lambda^{n(2b-2a)}$ , which always occurs as  $2b - 2a \leq 0$ . After such a flip,  $A_{t+1} = B_t$  while  $B_{t+1} = B_t$ . Thus  $\phi_{t+1} = 0$  and the change in distance between the two chains is  $-\phi_t = -\lambda^{n(2b-2a)}$ . The total contribution to the expected change in  $\phi(\mathcal{A}, \mathcal{B})$  from this move is  $-q \cdot \lambda^{n(2b-2a)}$ .

Similarly, the probability  $(x, y, d, o)$  selects edge  $f$  in  $B_t$  is also  $q = 1/(4n)$ , and these values do not yield a flippable edge in  $A_t$ . Edge  $f$  flips only if  $p < \lambda^{n(2a-2b)}$ , which occurs with probability  $\lambda^{n(2a-2b)} < 1$ . If this move occurs, then  $B_{t+1} = A_t = A_{t+1}$ , and the change in distance between  $A$  and  $B$  is again  $-\lambda^{n(2b-2a)}$ . The total contribution to the expected change in  $\phi(\mathcal{A}, \mathcal{B})$  from this move is

$$-q \cdot \lambda^{n(2a-2b)} \cdot \lambda^{n(2b-2a)} = -q.$$

While the two potential moves above decrease the distance between the coupled chains according to metric  $\phi$ , there are also moves that increase it. However, dyadic constraints limit the number of such moves, because two dyadic intervals of the same length cannot overlap nontrivially (Lemma 3.1). For  $A_t$ , this means the top and bottom edges of  $S$  are not flippable, because the vertical dimension of any rectangle resulting from such a flip could not be dyadic: it would be an interval of length  $2b$  overlapping  $S$ 's projection onto the  $y$ -axis, which is a dyadic interval of length  $2b$ . At first glance there are four other potential edge-flips for  $A_t$  involving  $S$ , specifically flips of the top and bottom halves of  $S$ 's left and right boundaries. However, again by Lemma 3.1, at most one of the left boundary and the right boundary of  $S$  contains flippable edges. Without loss of generality, we assume it is the right boundary of  $S$ , and label the two potentially flippable edges as  $g$  and  $h$ . Similarly, for  $B_t$ , at first glance there exist four other potential edge-flips involving  $S$ , specifically the left and right halves of  $S$ 's top and bottom boundaries. By Lemma 3.1, we assume without loss of generality that only portions of  $S$ 's bottom boundary are potentially flippable, and label the two potentially flippable edges as  $i$  and  $j$ .

Such edge-flips can only occur if  $A_t$  and  $B_t$  are tiled in the neighborhood of  $S$  as in Figure 3.6. To get an upper bound on  $\mathbb{E}[\Delta\phi_t]$ , we suppose this worst case neighborhood tiling exists. Edges  $g$  and  $h$  are each selected by values  $(x, y, d, o)$  in  $A_t$  with probability  $q$ ; both are then flipped with probability  $\lambda^{n(4a-b)}$ . The tiling  $A_{t+1}$  resulting from this flip is at distance  $\lambda^{n(b-4a)}$  from configuration  $A_t$ . The same selection  $(x, y, d, o)$  does not result in any flip in  $B_t$ , so  $B_{t+1} = B_t$ . The change in distance between  $\mathcal{A}$  and  $\mathcal{B}$  for each of these two moves is thus  $\lambda^{n(b-4a)}$ . In all, the contribution by these moves to the expected change in distance between the coupled chains is at most

$$2 \cdot q \lambda^{n(4a-b)} \cdot \lambda^{n(b-4a)} = 2q.$$

Similarly, edges  $i$  and  $j$  are selected to be flipped in  $B_t$  by values  $(x, y, d, o)$  with prob-




Figure 3.6: An area  $2/n$  rectangle  $S$  bisected by (a) horizontal edge  $e$  in  $A_t$  and (b) vertical edge  $f$  in  $B_t$ . Four “bad” edge-flips  $g, h, i, j$  exist only if  $A_t$  and  $B_t$  are tiled (up to reflection) in the neighborhood of  $S$  as shown.

ability  $q$ , and once selected, these edge-flips occur if  $p < \lambda^{n(4b-a)}$ , a bound which is at least 1 for  $a \geq 8b$ . The tiling  $B_{t+1}$  resulting from either flip is at distance  $\lambda^{n(4b-a)}$  from configuration  $B_t$ . The same values of  $(x, y, d, o)$  producing these moves also yield  $A_{t+1} = A_t$ . Thus the change in distance between  $\mathcal{A}$  and  $\mathcal{B}$  for these two moves is at most  $\lambda^{n(4b-a)}$ . In all, the contribution by these moves to the expected change in distance between the two chains in the coupling is at most  $2 \cdot q \cdot \lambda^{n(4b-a)}$ .

In total, we have shown

$$\begin{aligned}
\mathbb{E}[\phi_{t+1} - \phi_t] &\leq -q - q\lambda^{n(2b-2a)} + 2q + 2q\lambda^{n(4b-a)} \\
&= -q\lambda^{n(2b-2a)} (\lambda^{n(2a-2b)} + 1 - 2\lambda^{n(2a-2b)} - 2\lambda^{n(2b+a)}) \\
&= -q\phi_t (1 - \lambda^{n(2a-2b)} - 2\lambda^{n(2b+a)})
\end{aligned}$$



As  $a \geq 8b$  and  $a \geq 1/\sqrt{n}$ ,

$$2a - 2b \geq 2\left(a - \frac{1}{8}a\right) \geq a \geq 1/\sqrt{n}.$$

Additionally,  $2b + a \geq a \geq 1/\sqrt{n}$ . Thus,

$$\lambda^{n(2a-2b)} + 2\lambda^{n(2b+a)} \leq \lambda^{\sqrt{n}} + 2\lambda^{\sqrt{n}} = 3\lambda^{\sqrt{n}}$$

Provided  $\lambda < 3^{-1/\sqrt{n}}$ , as we assumed at the start of this section, we have that

$$\lambda^{n(2a-2b)} + 2\lambda^{n(2b+a)} \leq 3\lambda^{\sqrt{n}} < 1$$

Then, for some  $c > 0$ , where  $c = 1 - 3\lambda^{\sqrt{n}}$ , we have that

$$\mathbb{E}[\phi_{t+1} - \phi_t] \leq -q\phi_t \cdot c.$$

It follows that, as desired for the  $A_t$  and  $B_t$  we consider,

$$\mathbb{E}[\phi_{t+1}] \leq (1 - qc)\phi_t.$$

The closer  $\lambda$  is to the bound  $3^{-1/\sqrt{n}}$ , the smaller the value of  $c$  is.

**Case a = 2b.** The analysis of potential good moves and bad moves remains the same as the first case above, though certain probabilities and distances change. Initially,  $\phi(A_t, B_t) = \lambda^{n(2b-2a)} = \lambda^{-2nb}$ . We note that the contribution to the expected change in distance from good moves flipping edges  $e$  and  $f$  is still

$$-q(1 + \lambda^{n(2b-2a)}) = -q(1 + \lambda^{-2nb}).$$

The contributions to the expected change in distance from flipping edges  $g$  and  $h$  is still

$2q$ . However, for the edges  $i$  and  $j$ , once selected by  $(x, y, d, o)$ , flips now occur with probability  $q\lambda^{4b-a} = q\lambda^{2nb}$  rather than probability  $q$ . Such a move results in a change in distance between the chains in the coupling of  $\lambda^{n(a-4b)} = \lambda^{-2nb}$ . The expected contribution to the change in distance from these moves is now  $2q\lambda^{2nb}\lambda^{-2nb} = 2q$ .

In total, we see that in this case,

$$\begin{aligned}\mathbb{E}[\phi_{t+1} - \phi_t] &\leq -q(1 + \lambda^{-2nb}) + 4q \\ &= -q\lambda^{-2nb}(\lambda^{2nb} + 1 - 4\lambda^{2nb}) \\ &= -q\phi_t(1 - 3\lambda^{2nb}).\end{aligned}$$

Because we know that  $2ab = 1/n$ , we have  $a = 2b = 1/\sqrt{n}$ . This means  $3\lambda^{2nb} = 3\lambda^{\sqrt{n}}$ , and this value is less than one because  $\lambda < 3^{-1/\sqrt{n}}$ . Setting  $c = 1 - 3\lambda^{2nb} = 1 - 3\lambda^{\sqrt{n}} > 0$ , we have, as desired for these  $A_t$  and  $B_t$ ,

$$\mathbb{E}[\phi_{t+1}] \leq (1 - qc)\phi_t.$$

As above, the closer  $\lambda$  is to  $3^{-1/\sqrt{n}}$  the smaller constant  $c$  is. □

**Theorem 3.12.** *For any  $\lambda < 3^{-1/\sqrt{n}}$ , the edge-flip Markov chain  $\mathcal{M}_k^{edge}$  on dyadic tilings of size  $n = 2^k$  has mixing time at most  $O(n^2)$ .*

*Proof.* We apply the exponential metric theorem from [62] (Theorem 2.5), using the coupling  $(\mathcal{A}, \mathcal{B})$  and metric  $\phi$  defined above. Metric  $\phi$  satisfies the path requirement of Theorem 2.5 with  $U$  being the set of all pairs of tilings that are adjacent in  $\Omega_n$ , and by Lemma 3.10  $\phi$  takes on values in  $\{0\} \cup [1, S]$  for  $S = n \log(n)\lambda^{-n}$ . Additionally  $\mathcal{M}_k^{edge}$  is lazy (Lemma 3.7). We have also demonstrated (Lemma 3.11) that when  $A_t$  and  $B_t$  differ by a single flip there is a constant  $c$  such that  $\mathbb{E}[\phi_{t+1}] \leq (1 - qc)\phi_t$  whenever  $\lambda < 3^{-1/\sqrt{n}}$ .

By Theorem 2.5, recalling that  $q = 1/4n$  we conclude that

$$\begin{aligned} t_{\text{mix}}(\varepsilon) &\leq \frac{\ln(n \log n \lambda^{-n} \varepsilon^{-1})}{qc} = \frac{4n}{c} (\ln n + \ln \log n + n \ln \lambda^{-1} + \ln \varepsilon^{-1}) \\ &= O(n^2 + n \ln(\varepsilon^{-1})). \end{aligned}$$

When we assume  $\varepsilon = 1/4$ , as is standard practice, we see  $t_{\text{mix}} = t_{\text{mix}}(1/4) = O(n^2)$ .  $\square$

Now that we know the mixing time whenever  $\lambda < 3^{-1/\sqrt{n}}$ , it only remains to consider sufficiently large  $n$  to extend this result to all  $\lambda < 1$ .

**Theorem 3.13.** *For any constant  $\lambda < 1$ , the edge-flip Markov chain  $\mathcal{M}_k^{\text{edge}}$  on dyadic tilings of size  $n = 2^k$  has mixing time at most  $O(n^2)$ .*

*Proof.* For any constant  $0 < \lambda < 1$ , there is an  $n_0$  such that for all  $n > n_0$ ,  $3^{-1/\sqrt{n}} > \lambda$ . By Theorem 3.12, this implies that for all  $n > n_0$ , the mixing time of  $\mathcal{M}_k^{\text{edge}}$  is  $O(n^2)$ , which suffices to prove the claim.  $\square$

We briefly remark that this mixing time upper bound is within a factor of  $n/\log n$  of the true mixing time. We show this by presenting a lower bound on the mixing time of  $(n \log n)/4$ . This proof holds for Markov chain  $\mathcal{M}_k^{\text{edge}}$  and all values of  $\lambda$ , though better lower bounds when  $\lambda \geq 1$  will be presented in the next sections. We use the well-known result that half the diameter of the state space is a lower bound on mixing time (Theorem 2.7).

**Theorem 3.14.** *For any  $\lambda$ , the edge-flip Markov chain  $\mathcal{M}_k^{\text{edge}}$  on dyadic tilings of size  $n = 2^k$  has mixing time at least  $(n \log n)/4$ .*

*Proof.* Let  $\sigma_v$  denote the tiling consisting entirely of  $1/n \times 1$  vertical rectangles and let  $\sigma_h$  denote the tiling consisting entirely of  $1 \times 1/n$  horizontal rectangles. Starting at  $\sigma_v$ , the number of edge-flips required to move to  $\sigma_h$  is at least  $(n \log n)/2$ : one edge-flip at most doubles the width of two rectangles, and  $n$  rectangles must have their width doubled at least

$\log n$  times each to reach  $\sigma_h$ . This implies the diameter of  $\Omega_n$  under edge-flip moves is at least  $(n \log n)/2$ , which by Theorem 2.7 gives the claimed mixing time lower bound.  $\square$

We can also use Lemma 3.11 to give a relaxation time upper bound for  $\mathcal{M}_k^{edge}$ .

**Theorem 3.15.** *For any constant  $\lambda < 1$ , the edge-flip Markov chain  $\mathcal{M}_k^{edge}$  on dyadic tilings of size  $n = 2^k$  has relaxation time at most  $O(n)$ .*

*Proof.* We will use a coupling argument, applying Theorem 2.3 with the same coupling and distance metric  $\phi$  on dyadic tilings as above. To begin, we assume  $\lambda < 3^{-1/\sqrt{n}}$ .

Let  $A_t$  and  $B_t$  be the marginal dyadic tilings at some time  $t$  in the coupling; we do not assume that  $A_t$  and  $B_t$  differ by an edge flip. By the definition of this distance metric  $\phi$ , we know that there is some path from  $A_t$  to  $B_t$  in  $\Omega_k$  that includes, in order, states  $X_t^1, X_t^2, \dots, X_t^m$  for some  $m \geq 1$ , where

$$\phi(A_t, B_t) = \phi(A_t, X_t^1) + \sum_{i=1}^{m-1} \phi(X_t^i, X_t^{i+1}) + \phi(X_t^m, B_t)$$

For simplicity, we let  $A_t = X_t^0$  and  $B_t = X_t^{m+1}$ . Because all pairs  $X_t^i, X_t^{i+1}$  differ by an edge flip, if we suppose the coupling has these marginal tilings and let it evolve for one step, by Lemma 3.11 we know that, for  $q = 1/4n$  and  $c$  a constant, for each  $i = 0, 1, \dots, m$ ,

$$\mathbb{E}[\phi(X_{t+1}^i, X_{t+1}^{i+1}) \mid X_t^i, X_t^{i+1}] \leq (1 - qc)\phi(X_t^i, X_t^{i+1})$$

Because  $\phi(A_{t+1}, B_{t+1})$  is bounded above by the sum of distances along any path from  $A_{t+1}$  to  $B_{t+1}$ , in particular by the distances along the path  $X_{t+1}^1, X_{t+1}^2, \dots, X_{t+1}^m$ , then

$$\begin{aligned} \mathbb{E}[\phi(A_{t+1}, B_{t+1}) \mid A_t, B_t] &\leq \sum_{i=0}^m \mathbb{E}[\phi(X_{t+1}^i, X_{t+1}^{i+1}) \mid X_t^i, X_t^{i+1}] \\ &\leq (1 - qc) \sum_{i=0}^m \phi(X_t^i, X_t^{i+1}) \\ &\leq (1 - qc)\phi(A_t, B_t). \end{aligned}$$

Setting  $\beta = 1 - qc$ , we see that we have satisfied the hypotheses of the coupling theorem, Theorem 2.3. We conclude that the spectral gap  $\gamma$  of  $\mathcal{M}_k^{edge}$  satisfies

$$\gamma \geq 1 - \beta = qc = \frac{c}{4n}$$

We conclude the relaxation time of  $\mathcal{M}_k^{edge}$  for  $n = 2^k$  whenever  $\lambda < 3^{-1/\sqrt{n}}$  is at most  $O(n)$ .

For any constant  $0 < \lambda < 1$ , there is an  $n_0$  such that for all  $n > n_0$ ,  $3^{-1/\sqrt{n}} > \lambda$ . This implies that for all  $n > n_0$ , the relaxation time of  $\mathcal{M}_k^{edge}$  is  $O(n)$ , which suffices to prove the claim.  $\square$

This result will be used to show the behavior of  $\mathcal{M}_k^{edge}$  is different when  $\lambda < 1$  and when  $\lambda = 1$ : Section 3.5 establishes a super-linear lower bound on the relaxation time when  $\lambda = 1$ . The implications of this will be discussed more in Section 3.6.

### 3.4 Exponential Convergence when $\lambda > 1$

In contrast to the fast mixing results for  $\lambda < 1$  in the previous section, here we show that whenever  $\lambda > 1$  the mixing time of  $\mathcal{M}_k^{edge}$  is exponential in  $n = 2^k$ . Combined with the results of the previous section, this shows the existence of a phase transition for  $\mathcal{M}_k^{edge}$  at the critical point  $\lambda_c = 1$ . Sharp phase transitions such as these are conjectured for many systems, especially from statistical physics, but are often hard to rigorously confirm. Because of this it is remarkable that for dyadic tilings we were able to pin down the exact point at which the phase transition occurs.

We begin by rewriting  $|\sigma|$ , the sum of the lengths of all edges of  $\sigma$ , in way that will be more useful for our purposes. Let  $R$  be a rectangle of  $\sigma$ , and let  $w(R)$  be its width and  $l(R)$  be its length (height). All edges in  $\sigma$ , except for its left boundary and bottom boundary, are either part of the top boundary of some rectangle or the right boundary of some rectangle.

Using this logic, we can write

$$|\sigma| = 2 + \sum_{R \in \sigma} (l(R) + w(R)).$$

To show  $\mathcal{M}_k^{edge}$  takes exponential time to converge when  $\lambda > 1$ , we consider the tilings with at least one  $1 \times 1/n$  rectangle and those with at least one  $1/n \times 1$  rectangle. In order to go between these sets we must go through a tiling where all rectangles have width and length both at least  $2/n$  and thus each has length plus width at most  $1/2 + 2/n$ . We show these tilings are exponentially unlikely and thus our state space forms a bottleneck.

**Theorem 3.16.** *For any constant  $\lambda > 1$ , the edge-flip chain  $\mathcal{M}_k^{edge}$  on the set  $\Omega_k$  of dyadic tilings of size  $n = 2^k$  requires time  $\exp(\Omega(n^2))$  to mix.*

*Proof.* We first partition the state space into three sets:  $A$ , the set of tilings with at least one  $1/n \times 1$  vertical rectangle;  $B$ , the set of tilings with at least one  $1 \times 1/n$  horizontal rectangle; and  $C$ , the remainder, the set of tilings with no  $1/n \times 1$  or  $1 \times 1/n$  rectangles. We first calculate the stationary probabilities of each of these sets.

Notice that  $A$  contains the tiling  $\sigma_v$  where all rectangles are  $1/n \times 1$  and  $B$  contains the tiling  $\sigma_h$  where all rectangles are  $1 \times 1/n$ . Tiling  $\sigma_v$  consists of  $n$  rectangles, each with length plus width summing to  $1 + 1/n$ , as does  $\sigma_h$ , so we see that

$$|\sigma_v| = |\sigma_h| = n \left( 1 + \frac{1}{n} \right) + 2.$$

We conclude that

$$\begin{aligned} \pi(A) &\geq \pi(\sigma_v) = \frac{\lambda^{n^2(1+1/n)+2n}}{Z} = \frac{\lambda^{n^2+3n}}{Z}, \\ \pi(B) &\geq \pi(\sigma_h) = \frac{\lambda^{n^2(1+1/n)+2n}}{Z} = \frac{\lambda^{n^2+3n}}{Z}. \end{aligned}$$

For set  $C$ , every rectangle in every tiling has length and width both at least  $2/n$  and at most  $1/2$ . Each rectangle has length times width equal to 1, and we see that subject to this

constraint, length plus width is maximized at the extremal points when one is  $1/2$  and the other is  $2/n$ , that is, when  $l(R) + w(R) = 1/2 + 2/n$ . For any  $\sigma \in C$ , we see that

$$|\sigma| \leq n \left( \frac{1}{2} + \frac{2}{n} \right) + 2 = \frac{n}{2} + 4$$

As the number of tilings in set  $C$  is at most the total number of tilings in  $\Omega_k$ , and  $|\Omega_k| < 2^n$  by Lemma 3.3, we conclude that

$$\pi(C) \leq |C| \cdot \frac{\lambda^{n^2/2+4n}}{Z} \leq \frac{2^n \lambda^{n^2/2+4n}}{Z}$$

Using these bounds, we next bound the conductance of the Markov chain and then the mixing time using Theorem 2.8. We consider set  $S = A$ , and its complement  $\bar{S} = B \cup C$ . As  $\pi(A) = \pi(B)$ , because there is a length-preserving bijection between  $A$  and  $B$  given by rotations of 90 degrees, we conclude  $\pi(S) = \pi(A) < \pi(B \cup C) = \pi(\bar{S})$ , so in particular  $\pi(A) < 1/2$ . We see that, because there are no moves that transition from  $A$  directly into  $B$ , we must go through  $C$  first, implying

$$\Phi \leq \frac{1}{\pi(A)} \sum_{s_1 \in A, s_2 \in B \cup C} \pi(s_1) P(s_1, s_2) = \frac{1}{\pi(A)} \sum_{s_1 \in A, s_2 \in C} \pi(s_1) P(s_1, s_2).$$

Furthermore, using detailed balance and that fact that all entries in transition matrix  $P$  are at most one, we see that

$$\begin{aligned} \Phi &\leq \frac{1}{\pi(A)} \sum_{s_1 \in A, s_2 \in C} \pi(s_2) P(s_2, s_1) \\ &< \frac{1}{\pi(A)} \sum_{s_2 \in C} \pi(s_2) = \frac{\pi(C)}{\pi(A)} \leq \frac{2^n \lambda^{n^2/2+4n}}{\lambda^{n^2+3n}} = 2^n \lambda^{-n^2/2+n}. \end{aligned}$$

We conclude there exists a constant  $c_1 > 0$  such that for sufficiently large  $n$ ,  $\Phi < \lambda^{-c_1 n^2}$ .

Applying Theorem 2.8, which relates the conductance and mixing time of a Markov

chain, proves that for all  $\varepsilon > 0$ , the mixing time of  $\mathcal{M}_k^{edge}$  satisfies

$$t_{\text{mix}}(\varepsilon) \geq \left( \frac{\lambda^{c_1 n^2}}{4} - \frac{1}{2} \right) \log \left( \frac{1}{2\varepsilon} \right) = \Omega(\lambda^{c_1 n^2} \log \varepsilon^{-1}).$$

Letting  $\varepsilon = 1/4$  we have that  $t_{\text{mix}} = \Omega(\lambda^{c_1 n^2})$ , as desired.  $\square$

This concludes our proofs showing the existence of a phase transition at critical point  $\lambda_c = 1$  for Markov chain  $\mathcal{M}_k^{edge}$  on dyadic tilings of size  $n = 2^k$ . In the next section, we investigate the behavior of  $\mathcal{M}_k^{edge}$  at this critical point  $\lambda = 1$ .

### 3.5 Polynomial Convergence when $\lambda = 1$

The work in the previous two sections has characterized the mixing time of the biased edge-flip Markov chain  $\mathcal{M}_k^{edge}$  whenever  $\lambda \neq 1$ , as polynomial in  $n = 2^k$  when  $\lambda < 1$  and exponential in  $n$  when  $\lambda > 1$ . In this section we give a polynomial upper bound of  $O(n^{5.09})$  on the mixing time of  $\mathcal{M}_k^{edge}$  when  $\lambda = 1$ , resolving the open question of Janson, Randall, and Spencer [67] that had remained open from 2002 until we answered it in 2016. We also give a nontrivial lower bound on the mixing time when  $\lambda = 1$  of  $\Omega(n^{1.38})$ ; the implications of such a lower bound in statistical physics are discussed in 3.6.

#### 3.5.1 Proof Ideas

In addition to answering a long-standing open question and giving mixing time bounds for the last remaining value of  $\lambda$ , this result is also of more general interest because of its proof techniques. Building on similar work within the statistical physics community, this bisection/block moves approach has the potential to give (or improve) polynomial mixing time upper bounds for local Markov chains on other self-reducible structures. Using this technique to improve the mixing time upper bound for the edge-flip Markov chain for convex triangulations [87] is one focus of current work.



We identify a certain block structure on dyadic tilings that allows us to relate the spectral gap of the edge-flip Markov chain to that of another, simpler Markov chain. In the simpler Markov chain, which we refer to as the block dynamics, for each transition a large region of the tiling is selected and retiled uniformly at random, if possible. At the smallest scale,  $n = 4$ , these correspond to exactly the moves of the (lazy) edge-flip Markov chain. The structure of these block moves allows us to set up a recursion that relates the spectral gap of the edge-flip Markov chain for tilings of size  $n$  with that of sizes smaller than  $n$  and that of the block dynamics. This produces an inverse polynomial lower bound on the spectral gap of the edge-flip Markov chain.

Specifically, we adapt a bisection approach inspired by spin system analysis [86, 28]. We bound the spectral gap  $\gamma_k$  of the Markov chain  $\mathcal{M}_k^{edge}$  for dyadic tilings of size  $n = 2^k$  by the product of the spectral gap  $\gamma_{block}$  of the block dynamics Markov chain and the spectral gap  $\gamma_{k-1}$  of  $\mathcal{M}_{k-1}$ , and then use recursion to obtain

$$\gamma_k \geq \gamma_{k-1} \cdot \gamma_{block} \geq (\gamma_{block})^k = (\gamma_{block})^{\log n}.$$

To obtain a polynomial relaxation time and thus a polynomial mixing time, we complete our argument by showing that  $\gamma_{block}$  is constant using a coupling argument as in, e.g., Chapter 13 of [76]. The distance metric we use is a carefully weighted average of two different notions of distance between tilings. We do a case analysis and show this distance metric contracts by a factor of at least  $1 - 1/17$  in each step, implying the spectral gap  $\gamma_{block}$  is at least  $1/17$ . This gives a relaxation time of at most  $O(n^{\log 17}) = O(n^{4.09})$ , and a mixing time of at most  $O(n^{\log 17+1}) = O(n^{5.09})$ .

For our lower bound, we use a distinguishing statistic to show the mixing time and relaxation time of the edge-flip Markov chain  $\mathcal{M}_k^{edge}$  for dyadic tilings are at least  $\Omega(n^{1.38})$ ; again, see Chapter 13 of [76]. That is, we define a function  $f$  on the state space  $\Omega_k$  of all dyadic tilings of size  $n = 2^k$ . By considering the variance and Dirichlet form of  $f$ , and

using combinatorial properties of dyadic tilings, we can give an upper bound on the spectral gap and thus a lower bound on the relaxation and mixing times of  $\mathcal{M}_k^{edge}$ .

### 3.5.2 The Transition Matrix of $\mathcal{M}_k^{edge}$

In this section we will consider  $\mathcal{M}_k^{edge}$  as stated in Algorithm 1. As opposed to previous sections, we will be exploring spectral properties of  $\mathcal{M}_k^{edge}$ 's transition matrix  $P_{k,edge}$ . We briefly examine what  $P_{k,edge}$  looks like. For every valid edge flip of  $\mathcal{M}_k^{edge}$ , there are two choices of a rectangle  $R$  and an edge  $e$  that propose it. Since in each iteration  $\mathcal{M}_k^{edge}$  chooses each rectangle with probability  $1/n$  and one of its four sides with probability  $1/4$ , this implies every move between two tilings differing by an edge flip occurs with probability  $2/(4n) = 1/(2n) = 2^{-k-1}$ , so all nonzero off-diagonal entries of  $P_{k,edge}$  are  $2^{-k-1}$ . We let  $\gamma_k$  denote the spectral gap of  $\mathcal{M}_k^{edge}$ ; we omit the subscript  $edge$  from our notation  $\gamma_k$  for brevity. We will get an inverse polynomial lower bound on  $\gamma_k$ , which gives a polynomial upper bound on the relaxation and mixing times of  $\mathcal{M}_k^{edge}$ .

### 3.5.3 The Block Dynamics Markov Chain $\mathcal{M}_k^{block}$

To analyze the spectral gap of Markov chain  $\mathcal{M}_k^{edge}$ , we will appeal to a different Markov chain that uses larger block moves instead of single edge flips. While this Markov chain is not efficiently implementable because each step requires generating a large random dyadic tiling, it still plays an important role in our proofs. To define this new chain we depend critically on the bijection between tilings in  $\Omega_{k-1}$  and the left or right (resp. top or bottom) half of a tiling in  $\Omega_k$  that has a vertical (resp. horizontal) bisector, as discussed in the proof of Proposition 3.2. For  $k \geq 2$ , the block dynamics Markov chain  $\mathcal{M}_k^{block}$  on the state space  $\Omega_k$  of all dyadic tilings of size  $2^k$  is given by Algorithm 3.

Let  $P_{k,block}$  be the transition matrix of this Markov chain and let  $\gamma_{k,block}$  be its spectral gap. Any valid nonstationary transition of  $\mathcal{M}_k^{block}$  occurs with probability  $1/(4|\Omega_{k-1}|)$ . This Markov chain is not lazy, but it is aperiodic, irreducible, and reversible; this implies

---

**Algorithm 3** Block dynamics Markov chain  $\mathcal{M}_k^{block}$  on unbiased ( $\lambda = 1$ ) dyadic tilings of size  $n = 2^k$

---

- Beginning at any  $\sigma_0 \in \Omega_k$ , repeat:
- 1: Uniformly at random choose a tiling  $\rho \in \Omega_{k-1}$ .
  - 2: Uniformly at random choose *Left*, *Right*, *Top*, or *Bottom*.
  - 3: **if** *Left* was chosen and  $\sigma_i$  has a vertical bisector **then**
  - 4:     To obtain  $\sigma_{i+1}$ , retiling  $\sigma_i$ 's left half with  $\rho$ , under the mapping  $x \rightarrow x/2$ .
  - 5: **else if** *Right* was chosen and  $\sigma_i$  has a vertical bisector **then**
  - 6:     To obtain  $\sigma_{i+1}$ , retiling  $\sigma_i$ 's right half with  $\rho$ , under the mapping  $x \rightarrow (x + 1)/2$ .
  - 7: **else if** *Bottom* was chosen and  $\sigma_i$  has a horizontal bisector **then**
  - 8:     To obtain  $\sigma_{i+1}$ , retiling  $\sigma_i$ 's bottom half with  $\rho$ , under the mapping  $y \rightarrow y/2$ .
  - 9: **else if** *Top* was chosen and  $\sigma_i$  has a horizontal bisector **then**
  - 10:    To obtain  $\sigma_{i+1}$ , retiling  $\sigma_i$ 's top half with  $\rho$ , under the mapping  $y \rightarrow (y + 1)/2$ .
  - 11: **else**  $\sigma_{i+1} = \sigma_i$ .
- 

it is ergodic and thus has a unique stationary distribution which is uniform on  $\Omega_k$ . To prove these facts, we first recall from Section 3.1 the notion of *half-bisectors*. We say that a tiling  $x$  has a *left half-bisector* if the line segment from  $(0, 1/2)$  to  $(1/2, 1/2)$  does not intersect the interior of any dyadic rectangle. In an analogous way we can define a *right half-bisector* using the line segment from  $(1/2, 1/2)$  to  $(1, 1/2)$ , a *top half-bisector* using the line segment from  $(1/2, 1)$  to  $(1/2, 1/2)$ , and a *bottom half-bisector* using the line segment from  $(1/2, 1/2)$  to  $(1/2, 0)$ .

**Lemma 3.17.**  $\mathcal{M}_k^{block}$  is aperiodic, irreducible, and ergodic, and it converges to a unique stationary distribution that is the uniform distribution on  $\Omega_k$ .

*Proof.* We first show that  $\mathcal{M}_k^{block}$  is irreducible. For any tiling  $\sigma$ , to reach a tiling  $\tau \neq \sigma$  the following steps suffice. First, because  $\sigma$  must have a horizontal or vertical bisector, we can retiling the appropriate halves of  $\sigma$  to introduce any missing half-bisectors and reach a tiling with both a horizontal and a vertical bisector. Then, because  $\tau$  has at least one of these bisectors, we can retiling the appropriate halves as necessary to obtain  $\tau$ . Thus  $\mathcal{M}_k^{block}$  is irreducible.

To see  $\mathcal{M}_k^{block}$  is aperiodic, we note that for any  $\sigma$ ,  $P_{k,block}(\sigma, \sigma) \geq 1/(2|\Omega_{k-1}|)$ : each tiling  $\sigma$  has at least one bisector, and choosing one side of this bisector and the tiling of

$\Omega_{k-1}$  that is already present in that half of  $\sigma$  results in a stationary transition.

We conclude  $\mathcal{M}_k^{block}$  is ergodic, which means it converges to a unique stationary distribution over  $\Omega_k$ . For all  $\sigma \neq \tau$  we have that either

$$P_{k,block}(\sigma, \tau) = P_{k,block}(\tau, \sigma) = \frac{1}{4|\Omega_{k-1}|} \quad \text{or} \quad P_{k,block}(\sigma, \tau) = P_{k,block}(\tau, \sigma) = 0.$$

We conclude, via detailed balance, that the stationary distribution of  $\mathcal{M}_k^{block}$  is uniform over  $\Omega_k$ . □

We now bound the spectral gap of  $\mathcal{M}_k^{block}$  from below by a constant using coupling.

**Theorem 3.18.** *There exists a positive integer  $k_0$  such that for all  $k \geq k_0$ ,  $\gamma_{k,block} \geq 1/17$ .*

*Proof.* At a high level, we introduce a distance metric on dyadic tilings, and then give a coupling where the distance between two tilings decreases in expectation after one iteration by a multiplicative factor of at least  $1 - \frac{1}{17}$  for all  $k$  sufficiently large. Using Theorem 2.3, this implies the theorem.

We start defining the distance between two dyadic tilings  $x, y \in \Omega_k$ . For each of the four possible half-bisectors, let  $\ell_1$  be the number of such half-bisectors that are present in either  $x$  or  $y$ , but not in both of them. Also, for each of the four possible quadrants (top-left, top-right, bottom-left and bottom-right) of  $x$  and  $y$ , let  $\ell_2$  denote the number of such quadrants for which the rectangles in  $x$  intersecting that quadrant are not the same as the rectangles in  $y$  intersecting that quadrant. Then, introducing a parameter  $b > 0$  that we will take to be sufficiently large later, we define the distance between  $x$  and  $y$  as

$$d(x, y) = b\ell_1 + \ell_2.$$

For instance, consider the two dyadic tilings in Figures 3.1a and 3.1b. In this case we have  $\ell_1 = 1$  due to the left half-bisector that is present in (b) but not in (a), and  $\ell_2 = 3$  for top-left, top-right and bottom-left quadrants. The distance between these two tilings is then

$b + 3$ .

Our goal is to couple two instances of the block dynamics  $\mathcal{M}_k^{block}$ , one starting from a state  $A_t = x \in \Omega_k$  and the other from a state  $B_t = y \in \Omega_k$ , such that the distance between  $x$  and  $y$  contracts after one step of the chains. More precisely, letting  $\mathbb{E}_{x,y}$  denote the expectation with respect to the coupling, and if  $A_{t+1} = x'$  and  $B_{t+1} = y'$  are the dyadic tilings obtained after one step of each chain, respectively, we want to obtain a coupling and a value  $\Delta > 0$  such that

$$\mathbb{E}_{x,y}[d(x', y')] \leq (1 - \Delta)d(x, y) \quad \text{for all } x, y \in \Omega_k. \quad (3.1)$$

Once we have the above inequality, then Theorem 2.3 implies that  $\gamma_{k,block} \geq \Delta$ .

We will use the simple coupling between  $A_t = x$  and  $B_t = y$  that, in Steps 1 and 2 of Algorithm 3, chooses the same values for  $\rho$  and for the half to consider retiling, respectively. When we update the left (resp., right) half of  $x$  and  $\rho$  contains a horizontal bisector, note that  $x'$  will contain a left (resp., right) half-bisector. Similarly, if we update the top (resp., bottom) half of  $x$  and  $\rho$  contains a vertical bisector, then  $x'$  will contain a top (resp., bottom) half-bisector. In any of these cases, we say that the retiling yields a half-bisector of  $x$ . Recall that  $f_k := |\Omega_{k-2}|^2/|\Omega_{k-1}|$  is the number of ways to tile the left half of a tiling  $\sigma \in \Omega_k$  with a vertical bisector such that a left half-bisector is present;  $f_k$  has the same value and is defined similarly for the top, bottom, and right halves of dyadic tilings.

The remainder of the proof is devoted to showing that we can set  $b$  large enough so that (3.1) holds with  $\Delta = \frac{1}{17}$ . In order to see this, we will split into three cases, and show that (3.1) holds with  $\Delta = \frac{1}{17}$  for each case.

*Case 1:  $x$  and  $y$  have no common bisector.*

The maximum number of common half-bisectors of  $x$  and  $y$  in this case is two. Figure 3.7 illustrates the three possible configurations for the number of common half-bisectors of  $x$  and  $y$ . Consider first that  $x$  and  $y$  have no common half-bisector, which is illustrated in




Figure 3.7: Possible configurations for the half-bisectors of  $x$  and  $y$  in Case 1 of the proof of Theorem 3.18. In figure (c),  $i \in \{0, 1\}$  denotes how many grey quadrants are tiled identically in  $x$  and  $y$ .

Figure 3.7a and has  $d(x, y) = 4b + 4$ . Then, whichever half (left, right, top or bottom) is chosen to be retiled, note that either  $x$  or  $y$  is actually retiled, but never both. With probability  $\frac{|\Omega_{k-2}^2|}{|\Omega_{k-1}|} = f_k$  the retiling yields a half-bisector, which increases the number of common half-bisectors between  $x$  and  $y$ , and thus decreases their distance by  $b$ . Hence, using that  $f_k \geq 1/2$ , we have

$$\mathbb{E}_{x,y}[d(x', y')] = d(x, y) - f_k b \leq 4b + 4 - \frac{b}{2} < \left(1 - \frac{1}{17}\right) (4b + 4),$$

where the last step is true by setting  $b$  large enough (in this case,  $b \geq 1$  suffices).

Now consider that  $x$  and  $y$  have one common half-bisector, and use Figure 3.7b as a reference, with  $x$  being the left tiling and  $y$  being the right tiling. We have  $d(x, y) = 3b + 4$ . If we retiling the left or right halves, so only  $x$  gets retiled, and the retiling yields a half-bisector, then the number of common half-bisectors of  $x$  and  $y$  decreases by 1. A similar behavior happens if we retiling the top half. However, if we retiling the bottom half, and the retiling does not yield a half-bisector, then the number of common half-bisectors decreases by 1. Hence, using that  $f_k \geq 1/2$ , we obtain

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{3f_k b}{4} + \frac{(1 - f_k)b}{4} \leq 3b + 4 - \frac{b}{4} < \left(1 - \frac{1}{17}\right) (3b + 4),$$

where the last step is true by setting  $b$  large enough (in this case,  $b \geq 4$  suffices).

Finally, suppose  $x$  and  $y$  have two common half-bisectors, as illustrated in Figure 3.7c,

where they may or may not be tiled the same in the quadrant bounded by these common half-bisectors. In this case  $d(x, y) = 2b + 4 - i$ , where  $i = 1$  if they agree on this quadrant and  $i = 0$  otherwise. Retiling the left and top halves can yield a new common half-bisector, while retiling the right and bottom halves may remove a common half-bisector. Moreover, if  $i = 1$  and we retiling the right or bottom halves, the tilings of the bottom-right quadrant of  $x$  and of  $y$  may become different, increasing the distance between  $x$  and  $y$  by 1. Putting these together, we have

$$\begin{aligned} \mathbb{E}_{x,y}[d(x', y')] &\leq d(x, y) - \frac{2f_k b}{4} + \frac{2(1 - f_k)b}{4} + i\frac{2}{4} \\ &\leq 2b + 4 - \frac{i}{2} - \frac{(2f_k - 1)b}{2} \\ &= \frac{(5 - 2f_k)b}{2} + 4 - \frac{i}{2}. \end{aligned}$$

Since  $f_k \rightarrow \frac{\sqrt{5}-1}{2}$  as  $k \rightarrow \infty$ , the right-hand side above goes to  $\left(\frac{6-\sqrt{5}}{2}\right)b + 4 - \frac{i}{2}$ . In particular, for  $k \geq 10$ , the coefficient of  $b$  above satisfies  $\frac{5-2f_k}{2} < 2\left(1 - \frac{1}{17}\right)$ , and so we can set  $b$  large enough so that  $\mathbb{E}_{x,y}[d(x', y')] \leq \left(1 - \frac{1}{17}\right)(2b + 4 - i)$ . We note this is the tight case, as  $\frac{6-\sqrt{5}}{2} > 2\left(1 - \frac{1}{16}\right)$ , so this particular coupling and distance metric cannot be used to show the spectral gap is at least  $1/16$ . This concludes the first case.

*Case 2:  $x$  and  $y$  have a common bisector, but neither  $x$  nor  $y$  has both bisectors.*

Without loss of generality we assume  $x$  and  $y$  both have a vertical bisector and neither has a horizontal bisector. Each of  $x$  and  $y$  has at least 2 and at most 3 half-bisectors. Figure 3.8 illustrates the four possible configurations for the number of half-bisectors of  $x$  and  $y$ ; the shaded quadrants are those where  $x$  and  $y$  could have the same tiling. In all the situations of Figure 3.8, if we retiling the left or right halves, then we match up the configuration of  $x$  and  $y$  in that half. In particular, if  $x$  and  $y$  don't agree on the presence of a left half-bisector, then they also do not have the same tiling of the top left or bottom left quadrants, so the decrease in distance due to a retiling of the left half, a move that occurs with probability  $1/4$ , is  $(b + 2)$ . If  $x$  and  $y$  agree on the presence of a left half-bisector and have the same




Figure 3.8: Possible configurations for the half-bisectors of  $x$  and  $y$  in Case 2 of the proof of Theorem 3.18. The value of  $i \in \{0, 1, 2, 3\}$  denotes the number of grey quadrants which are tiled identically in  $x$  and  $y$ .




Figure 3.9: Possible configurations for the half-bisectors of  $x$  and  $y$  in Case 3 of the proof of Theorem 3.18. The value of  $i \in \{0, 1, 2, 3\}$  denotes the number of grey quadrants which are tiled identically in  $x$  and  $y$ .

tiling on  $i' \in \{0, 1, 2\}$  of the two left quadrants, then the decrease in distance due to a retiling of the left half is  $(2 - i')$ . The same holds for right half-bisectors and retilings of the right half. As there are no moves of the coupling that can increase the distance between  $x$  and  $y$ , it can be shown that in all of the cases shown in Figure 3.8 the distance decreases by  $1/4$  in expectation. Hence,

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{d(x, y)}{4} \leq \left(1 - \frac{1}{17}\right) d(x, y),$$

which concludes the second case.

*Case 3:  $y$  has both vertical and horizontal bisectors.*

Here there are three situations, depending on whether  $x$  has two, three or four half-bisectors; see Figure 3.9. In the situation of Figure 3.9a, if the left or right halves are retiled, then we match up  $x$  and  $y$  in that half, decreasing the distance by  $b + 2$ . But if we retiling the top or bottom halves, then we may increase the distance by  $b$  if the retiling does not yield a



half-bisector. Hence,

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{2(b+2)}{4} + \frac{2(1-f_k)b}{4} = \frac{(4-f_k)b}{2} + 3.$$

Since  $\frac{4-f_k}{2} \rightarrow \frac{9-\sqrt{5}}{4} < (1 - \frac{1}{17})$ , the right-hand side above is smaller than  $(1 - \frac{1}{17})(2b+4)$  when  $k$  and  $b$  are large enough. A similar situation occurs in Figure 3.9b, but the distance increases a bit more when the top or bottom half is retiled as quadrants that were equal in  $x$  and  $y$  may become different. In this case, we have

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{(b+4-i)}{4} + \frac{2(1-f_k)b}{4} + \frac{2}{4} = \frac{(5-2f_k)b}{4} + \frac{6-i}{4}.$$

Since  $\frac{5-2f_k}{4} \rightarrow \frac{6-\sqrt{5}}{4} < (1 - \frac{1}{17})$ , the right-hand side above is smaller than the value of  $(1 - \frac{1}{17})(b+4-i)$  when  $k$  and  $b$  are large enough; this is the second tight case, where we see contraction by a factor of  $1 - \frac{1}{17}$  but not by  $1 - \frac{1}{16}$ . Finally, for the situation in Figure 3.9c, regardless of which half we choose to retiling, the distance will not increase; if we choose a half containing a quadrant on which  $x$  and  $y$  differ, the distance will decrease. Each quadrant on which  $x$  and  $y$  differ is contained in two halves and thus is retiled so that  $x$  and  $y$  agree there with probability  $1/2$ . That is,

$$\mathbb{E}_{x,y}[d(x', y')] \leq d(x, y) - \frac{d(x, y)}{2} \leq \left(1 - \frac{1}{17}\right) d(x, y).$$

This concludes the third case. We have shown that for all possible tilings  $x$  and  $y$ , it holds that  $\mathbb{E}_{x,y}[d(x', y')] \leq (1 - \frac{1}{17}) d(x, y)$ . By Theorem **[thm:coupling]**, this implies  $\gamma_{k,block} \geq \frac{1}{17}$  for all  $k$  sufficiently large, as desired.  $\square$

This concludes our proof that the spectral gap  $\gamma_{k,block}$  of  $\mathcal{M}_k^{block}$  is bounded below by a constant. In the next section we used this result to give a lower bound that is an inverse polynomial in  $n = 2^k$  on  $\gamma_k$ , the spectral gap of  $\mathcal{M}_k^{edge}$ , which yields polynomial upper bounds on the mixing and relaxation times of this chain.

### 3.5.4 A Polynomial Upper Bound on the Mixing and Relaxation Times of $\mathcal{M}_k^{edge}$

Recall we wish to show the mixing time of  $\mathcal{M}_k^{edge}$  is polynomial in  $n = 2^k$ , not polynomial in  $k$ . We show the spectral gap  $\gamma_k$  of  $\mathcal{M}_k^{edge}$  and the spectral gap  $\gamma_{k-1}$  of  $\mathcal{M}_{k-1}^{edge}$  differ by a multiplicative constant (specifically,  $1/17$ ) by appealing to the Dirichlet forms of both of these Markov chains as well as the block dynamics Markov chain  $\mathcal{M}_k^{block}$ . We can then use recursion to show  $\gamma_k$  is bounded below by  $(1/17)^k$ , which, because  $k = \log n$ , gives a polynomial upper bound on the relaxation time and thus on the mixing time of  $\mathcal{M}_k^{edge}$ .

For any function  $f : \Omega_k \rightarrow \mathbb{R}$ , we will denote the Dirichlet form of  $f$  with respect to transition matrix  $P_{k,edge}$  and the uniform stationary distribution as  $\mathcal{E}_{k,edge}(f)$ . The Dirichlet form of  $f$  with respect to transition matrix  $P_{k,block}$  and the uniform stationary distribution will be  $\mathcal{E}_{k,block}(f)$ . We will let the variance of function  $f$  on  $\Omega_k$  with respect to the uniform stationary distribution be  $\text{var}_k(f)$ . Here the  $k$  indicates which state space  $\Omega_k$  we are considering, rather than which distribution on  $\Omega_k$  the variance is taken with respect to; all variances we consider will be with respect to the uniform distribution.

Because we consider two different Markov chains on the same state space  $\Omega_k$ , there are two different notions of adjacencies on this state space, each corresponding to the moves of one of these Markov chains. For  $x, y \in \Omega_k$ , we say  $x \sim_e y$  if  $x$  and  $y$  differ by a single edge flip move of  $\mathcal{M}_k^{edge}$  and  $x \sim_b y$  if  $x$  and  $y$  differ by a single move of the block dynamics chain  $\mathcal{M}_k^{block}$ . More specifically, if  $x$  and  $y$  differ by a retiling of their left half (implying  $x$  and  $y$  both have a vertical bisector and are the same on their right half), we say  $x \sim_L y$ ; then  $x \sim_R y$ ,  $x \sim_T y$ , and  $x \sim_B y$  are defined similarly for the right, top, and bottom halves.

**Theorem 3.19.** *For any  $k \geq 2$ , the spectral gap  $\gamma_k$  of the edge-flip Markov chain  $\mathcal{M}_k^{edge}$  satisfies*

$$\gamma_k \geq \gamma_{k,block} \cdot \gamma_{k-1}$$

*Proof.* We begin by relating the Dirichlet forms for block dynamics and for the edge-flip dynamics, which will allow comparison of their spectral gaps. Recall that for any function  $f : \Omega_k \rightarrow \mathbb{R}$ ,

$$\mathcal{E}_{k,block}(f) = \frac{1}{2} \sum_{x \sim_b y \in \Omega_k} \pi(x) P_{k,block}(x, y) (f(x) - f(y))^2.$$

This sum can be split into four terms, corresponding to the type of block move (left, right, top, or bottom) transforming  $x$  into  $y$ . If  $x$  and  $y$  differ only in their top-left quadrants, then  $x$  could transition to  $y$  via either a left block move or a top block move; each of these moves occurs with probability  $\frac{1}{4|\Omega_{k-1}|}$ , and the total probability of  $P_{k,block}(x, y) = \frac{1}{2|\Omega_{k-1}|}$  will be split correspondingly between the terms for left block moves and top block moves.

We now analyze the first of these terms, containing all  $x, y$  differing by a retiling of their left halves. For  $x_L, x_R \in \Omega_{k-1}$ , by  $x_L x_R$  below we mean the tiling in  $\Omega_k$  with a vertical bisector whose left half is  $x_L$  under the map  $x \rightarrow x/2$  and whose right half is  $x_R$  under the map  $x \rightarrow (x+1)/2$ .

$$\begin{aligned} \mathcal{E}_{k,block}^L &= \frac{1}{2} \sum_{x \sim_L y} \frac{1}{|\Omega_k|} \frac{1}{4|\Omega_{k-1}|} (f(x) - f(y))^2 \\ &= \frac{1}{8} \sum_{x_R \in \Omega_{k-1}} \sum_{x_L, y_L \in \Omega_{k-1}} \frac{1}{|\Omega_k|} \frac{1}{|\Omega_{k-1}|} (f(x_L x_R) - f(y_L x_R))^2 \\ &= \frac{1}{4} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \left( \frac{1}{2} \sum_{x_L, y_L \in \Omega_{k-1}} \frac{1}{|\Omega_{k-1}|^2} (f(x_L x_R) - f(y_L x_R))^2 \right). \end{aligned}$$

We note that the second sum above is over all pairs of tilings in  $\Omega_{k-1}$ . While the Dirichlet form of a function sums over all pairs of states that differ by a transition of a Markov chain, the variance of a function sums over all pairs of states, regardless of the local structure imposed on the state space by the Markov chain. In fact, we have written the second sum above suggestively, and note that it is in fact a variance of a function over the state space  $\Omega_{k-1}$ . For each  $x_R \in \Omega_{k-1}$ , the function  $f|_{x_R} : \Omega_{k-1} \rightarrow \mathbb{R}$  given by  $f|_{x_R}(z) = f(zx_R)$

has variance  $\text{var}_{k-1}(f|_{x_R})$  (with respect to the uniform distribution) that is exactly equal to the term in parentheses above. Because the variance of a function is the same regardless of which transitions on the state space we are considering, it is through this variance we can relate  $\mathcal{E}_{k,block}$ , which we have calculated above, to a Dirichlet form for edge-flip moves. That is, by Proposition 2.2, we can bound this variance with the Dirichlet form of  $f|_{x_R}$  associated to  $P_{k-1,edge}$  and the spectral gap  $\gamma_{k-1}$  of  $\mathcal{M}_{k-1}$ . Thus,

$$\mathcal{E}_{k,block}^L = \frac{1}{4} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \text{var}_{k-1}(f|_{x_R}) \leq \frac{1}{4} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \frac{\mathcal{E}_{k-1,edge}(f|_{x_R})}{\gamma_{k-1}}$$

We now see that the Dirichlet form for the edge-flip Markov chain on  $\Omega_{k-1}$  is

$$\begin{aligned} \mathcal{E}_{k-1,edge}(f|_{x_R}) &= \frac{1}{2} \sum_{\substack{x_L, y_L \in \Omega_{k-1} \\ x_L \sim_e y_L}} \pi(x_L) P(x_L, y_L) (f(x_L x_R) - f(y_L x_R))^2 \\ &= \sum_{\substack{x_L, y_L \in \Omega_{k-1} \\ x_L \sim_e y_L}} \frac{1}{|\Omega_{k-1}|} \frac{1}{2n} (f(x_L x_R) - f(y_L x_R))^2 \end{aligned}$$

Using this expression, we see that

$$\begin{aligned} \mathcal{E}_{k,block}^L(f) &\leq \frac{1}{4\gamma_{k-1}} \sum_{x_R \in \Omega_{k-1}} \frac{|\Omega_{k-1}|}{|\Omega_k|} \left( \sum_{\substack{x_L, y_L \in \Omega_{k-1} \\ x_L \sim_e y_L}} \frac{1}{|\Omega_{k-1}|} \frac{1}{2n} (f(x_L x_R) - f(y_L x_R))^2 \right) \\ &= \frac{1}{4\gamma_{k-1}} \sum_{\substack{x, y \in \Omega_k \\ x \sim_e y \\ x \sim_L y}} \frac{1}{|\Omega_k|} \frac{1}{2n} (f(x) - f(y))^2. \end{aligned}$$

We now compare this to the Dirichlet form for the edge flip Markov chain on  $\Omega_k$ , which we recall is

$$\mathcal{E}_{k,edge}(f) = \frac{1}{2} \sum_{\substack{x, y \in \Omega_k \\ x \sim_e y}} \frac{1}{|\Omega_k|} \frac{1}{2n} (f(x) - f(y))^2.$$

We note for every  $x, y \in \Omega_k$  such that  $x \sim_e y$ , at least one of and at most two of  $x \sim_L y$ ,

$x \sim_R y$ ,  $x \sim_T y$ , and  $x \sim_B y$  hold. Thus each summand of  $\mathcal{E}_{k,edge}(f)$  appears at most twice as a summand of

$$\mathcal{E}_{k,block}(f) = \mathcal{E}_{k,block}^L(f) + \mathcal{E}_{k,block}^R(f) + \mathcal{E}_{k,block}^T(f) + \mathcal{E}_{k,block}^B(f).$$

It follows that

$$\mathcal{E}_{k,block}(f) \leq \frac{1}{4\gamma_{k-1}} \cdot 2 \cdot (2\mathcal{E}_{k,edge}(f)) = \frac{\mathcal{E}_{k,edge}(f)}{\gamma_{k-1}}.$$

This implies that for any  $f$ , first by Proposition 2.2 and then using the above equation, that

$$\text{var}_k(f) \leq \frac{\mathcal{E}_{k,block}(f)}{\gamma_{k,block}} \leq \frac{\mathcal{E}_{k,edge}(f)}{\gamma_{k,block} \cdot \gamma_{k-1}}.$$

Let  $f$  be chosen to be the function achieving equality in  $\text{var}_k(f) \leq \frac{\mathcal{E}_{k,edge}(f)}{\gamma_{k,block} \cdot \gamma_{k-1}}$ . We conclude

$$\gamma_k = \frac{\mathcal{E}_{k,edge}(f)}{\text{var}_k(f)} \geq \gamma_{k,block} \cdot \gamma_{k-1}.$$

□

Recall Theorem 3.18 states that  $\gamma_{k,block}$  is at least  $1/17$  for sufficiently large  $k$ . This can be used to bound the spectral gap, the relaxation time, and finally the mixing time of  $\mathcal{M}_k^{edge}$ .

**Theorem 3.20.** *The relaxation time of the edge-flip Markov chain for dyadic tilings of size  $n$  is at most  $O(n^{\log 17})$ . As a consequence, the mixing time of this chain is at most  $O(n^{1+\log 17})$ .*

*Proof.* By Theorems 3.18 and 3.19, the spectral gap of  $\mathcal{M}_k^{edge}$  satisfies

$$\gamma_k \geq \frac{1}{17} \gamma_{k-1} \geq 17^{-(k-k_0)} \gamma_{k_0},$$

where  $k_0$  is the constant value from Theorem 3.18. Since  $\gamma_{k_0}$  is a constant that does not

depend on  $n$ ,

$$\gamma_k = \Omega(17^{-k}) = \Omega(n^{-\log 17}) = \Omega(n^{-4.09}).$$

Because  $\mathcal{M}_k^{edge}$  is a lazy Markov chain (Lemma 3.7), its relaxation time satisfies

$$t_{\text{rel}} = O(n^{\log 17}) = O(n^{4.09}).$$

To use this to bound the mixing time of  $\mathcal{M}_k^{edge}$ , we appeal to Proposition 2.1, though we first must calculate  $\pi_{\min}$ . For  $\pi$  the uniform distribution,  $\min_{x \in \Omega_k} \pi(x) = 1/|\Omega_k|$ . By Lemma 3.3, the number of dyadic tilings of size  $n = 2^k$  satisfies  $|\Omega_k| < 2^n$ , implying  $1/\pi_{\min} = |\Omega_k| < 2^n$ . We conclude

$$t_{\text{mix}} = O(n^{1+\log 17}) = O(n^{5.09}).$$

□

This result resolves the open question of [67] from 2002 and gives the first polynomial upper bound on the mixing time of  $\mathcal{M}_k^{edge}$ .

### 3.5.5 A Nontrivial Lower Bound on Convergence when $\lambda = 1$

In this section we provide lower bounds on the mixing and relaxation times of  $\mathcal{M}_k^{edge}$  when  $\lambda = 1$  that complement the upper bounds given in the previous section. The trivial lower bound for the mixing time is  $\Omega(n \log n)$ , which, as proved in Theorem 3.14, is a simple consequence of the fact that the diameter of the Markov chain is of order  $n \log n$ . In the theorem below we improve this bound to  $\Omega(n^{1.38})$  for both relaxation and mixing times. The implications of this result and what it says about statistical physics are then discussed in Section 3.6.

To begin, we define the following subsets of  $\Omega_k$ :

$$\Omega_k^+ = \{x \in \Omega_k : x \text{ has both a horizontal and a vertical bisector}\},$$

$$\Omega_k^\perp = \{x \in \Omega_k : x \text{ has a vertical bisector}\}, \text{ and}$$

$$\Omega_k^- = \{x \in \Omega_k : x \text{ has a horizontal bisector}\}.$$

By definition, we have  $\Omega_k^+ = \Omega_k^\perp \cap \Omega_k^-$ . We start with the following simple lemma.

**Lemma 3.21.** *For all  $k \geq 2$ , we have*

$$\frac{|\Omega_k|}{|\Omega_k^+|} = \frac{2}{f_k^2} - 1 \geq 2\phi + 1.$$

Furthermore,  $\lim_{k \rightarrow \infty} \frac{|\Omega_k|}{|\Omega_k^+|} = 2\phi + 1$ , where  $\phi = \frac{\sqrt{5}+1}{2}$  is the golden ratio.

*Proof.* Using that  $|\Omega_k^+| = |\Omega_{k-2}|^4$ , and Proposition 3.2, we have

$$\frac{|\Omega_k|}{|\Omega_k^+|} = \frac{2|\Omega_{k-1}|^2 - |\Omega_{k-2}|^4}{|\Omega_{k-2}|^4} = \frac{2}{f_k^2} - 1.$$

By Lemma 3.5,  $f_k \leq \frac{\sqrt{5}-1}{2} = \frac{1}{\phi}$ , and using the identity  $\phi^2 = 1 + \phi$  we see that

$$\frac{|\Omega_k|}{|\Omega_k^+|} = \frac{2}{f_k^2} - 1 \geq 2\phi^2 - 1 = 2 + 2\phi - 1 = 2\phi + 1.$$

Furthermore, because  $\lim_{k \rightarrow \infty} f_k = 1/\phi$ , we conclude

$$\lim_{k \rightarrow \infty} \frac{|\Omega_k|}{|\Omega_k^+|} = \lim_{k \rightarrow \infty} \frac{2}{f_k^2} - 1 = \frac{2}{(\lim_{k \rightarrow \infty} f_k)^2} - 1 = \frac{2}{(1/\phi)^2} - 1 = 2\phi - 1.$$

□

We will also require the following technical estimate.




Figure 3.10: The construction of a tiling to count  $\prod_{i=0}^{k-2} |\Omega_i|^2$ . A rectangle with number  $a$  indicates that we tile it with a tiling from  $\Omega_{k-a}$ .

**Lemma 3.22.** *For any  $k \geq 2$ , we have*

$$\frac{1}{|\Omega_k|} \prod_{i=0}^{k-2} |\Omega_i|^2 \leq \phi^{-2k+2}$$

*Proof.* We will show how to estimate  $\prod_{i=0}^{k-2} |\Omega_i|^2$  via the construction of a tiling in  $\Omega_k$ . We start with a tiling with both a horizontal and a vertical bisector, as in Figure 3.10(a). Then we inductively do the following. Both quadrants of the left half are tiled independently with a uniformly random tiling from  $\Omega_{k-2}$ . In the top-right quadrant, we add a vertical bisector and complete the two halves of this quadrant with independent, uniformly random tilings from  $\Omega_{k-3}$ . Finally, in the bottom-right quadrant, we create a horizontal and a vertical bisector, reaching the tiling in Figure 3.10(b). Then we take this bottom-right quadrant, and iterate the procedure above; see Figure 3.10(c,d) for the configurations after one and two more iterations. This iteration continues until creating a bisector will result in rectangles of area less than  $2^{-k}$ . In the case where an attempt is made to divide a rectangle of area  $2^{-k+1}$  into four rectangles of equal area by adding both a horizontal and vertical bisector, we instead add just a horizontal bisector, resulting in two rectangles each of area  $2^{-k}$ .

Let  $\Upsilon_k \subset \Omega_k$  be the set of tilings obtained in this way. Note that the number of tilings in  $\Upsilon_k$  is exactly  $\prod_{i=0}^{k-2} |\Omega_i|^2$ . Since  $\Upsilon_k \subset \Omega_k^+$ , we have that  $\frac{|\Upsilon_k|}{|\Omega_k|} \leq \frac{|\Omega_k^+|}{|\Omega_k|}$ , where the first expression is exactly the value we wish to bound. Using the construction above until Fig-



ure 3.10(b), we obtain that

$$\frac{|\Upsilon_k|}{|\Omega_k|} \leq \frac{|\Omega_k^+|}{|\Omega_k|} \frac{|\Omega_{k-2}^|}{|\Omega_{k-2}|},$$

where the second factor stands for the fact that the top-right quadrant must contain a vertical bisector. Iterating this in the bottom-right quadrant, we obtain

$$\frac{|\Upsilon_k|}{|\Omega_k|} \leq \frac{|\Omega_k^+|}{|\Omega_k|} \frac{|\Omega_{k-2}^|}{|\Omega_{k-2}|} \frac{|\Omega_{k-2}^+|}{|\Omega_{k-2}|} \frac{|\Omega_{k-4}^|}{|\Omega_{k-4}|} \dots \quad (3.2)$$

Proposition 3.2 gives that

$$\frac{|\Omega_k^|}{|\Omega_k|} = \frac{|\Omega_k| + |\Omega_{k-2}|^4}{2|\Omega_k|} = \frac{1}{2} \left( 1 + \frac{|\Omega_{k-2}^+|}{|\Omega_k|} \right) \leq \frac{1}{2} \left( 1 + \frac{1}{2\phi + 1} \right) = \frac{\phi^2}{2\phi + 1},$$

where the inequality follows from Lemma 3.21. For even  $k$ , because  $|\Omega_0^| = 0$  the last term we can obtain in (3.2) is  $\frac{|\Omega_2^+|}{|\Omega_2|}$ , so we can write

$$\begin{aligned} \frac{|\Upsilon_k|}{|\Omega_k|} &\leq \left( \prod_{i=0}^{k/2-2} \frac{|\Omega_{k-2i}^+|}{|\Omega_{k-2i}|} \cdot \frac{|\Omega_{k-2i-2}^|}{|\Omega_{k-2i-2}|} \right) \frac{|\Omega_2^+|}{|\Omega_2|} \\ &\leq \frac{1}{2\phi + 1} \left( \frac{1}{2\phi + 1} \cdot \frac{\phi^2}{2\phi + 1} \right)^{\frac{k}{2}-1} \\ &= \frac{\phi^{-2k+4}}{2\phi + 1} \\ &\leq \phi^{-2k+2}, \end{aligned}$$

where the last two expressions come from, respectively, identities for  $\phi$  and the easily-checked inequality  $2\phi + 1 > \phi^2$ . When  $k$  is odd, the last term in (3.2) is  $\frac{|\Omega_1^|}{|\Omega_1|}$  because  $|\Omega_1^+| = 0$ , so we can write

$$\frac{|\Upsilon_k|}{|\Omega_k|} \leq \left( \prod_{i=0}^{(k-3)/2} \frac{|\Omega_{k-2i}^+|}{|\Omega_{k-2i}|} \cdot \frac{|\Omega_{k-2i-2}^|}{|\Omega_{k-2i-2}|} \right) \leq \left( \frac{1}{2\phi + 1} \cdot \frac{\phi^2}{2\phi + 1} \right)^{\frac{k-1}{2}} \leq \phi^{-2k+2},$$

where again the last expression is the result of applying identities for  $\phi$  and simplifying.  $\square$

To prove our lower bound result, we will use a distinguishing statistic (Section 2.3.3). In particular, as our distinguishing statistic we consider the function  $f : \Omega_k \rightarrow \{0, 1\}$  such that

$$f(x) = \begin{cases} 1 & \text{if } x \in \Omega_k^| \\ 0 & \text{otherwise} \end{cases} \quad (3.3)$$

For a function  $f$  on  $\Omega_k$ , recall that  $\text{var}_k f$  denotes the variance of  $f$  with respect to the uniform measure on  $\Omega_k$ .

**Lemma 3.23.** *With  $f : \Omega_k \rightarrow \{0, 1\}$  as in (3.3), we have that*

$$\lim_{k \rightarrow \infty} \text{var}_k(f) = \sqrt{5} - 2.$$

*Proof.* We start by applying the definition of the variance of a function taken with respect to the uniform distribution  $\pi(x) = 1/|\Omega_k|$ .

$$\begin{aligned} \text{var}_k(f) &= \sum_{x, y \in \Omega_k} \pi(x)\pi(y)(f(x) - f(y))^2 \\ &= \sum_{x, y \in \Omega_k} \frac{1}{|\Omega_k|} \cdot \frac{1}{|\Omega_k|} \cdot (f(x) - f(y))^2 \\ &= \sum_{x \in \Omega_k^|} \sum_{y \in \Omega_k \setminus \Omega_k^|} \frac{1}{|\Omega_k|^2} \\ &= \frac{|\Omega_k^| \cdot |\Omega_k \setminus \Omega_k^|}{|\Omega_k|^2}. \end{aligned}$$

Since  $|\Omega_k^| = |\Omega_{k-1}|^2$ , using Proposition 3.2 we obtain

$$|\Omega_k^| = \frac{|\Omega_k| + |\Omega_{k-2}|^4}{2} = \frac{|\Omega_k| + |\Omega_k^+|}{2}, \quad (3.4)$$

and

$$|\Omega_k \setminus \Omega_k^| = |\Omega_k| - |\Omega_k^| = \frac{|\Omega_k| - |\Omega_k^+|}{2}. \quad (3.5)$$

Plugging in (3.4) and (3.5) to our above expression for the variance, we see that

$$\text{var}_k(f) = \frac{|\Omega_k^+| \cdot |\Omega_k \setminus \Omega_k^+|}{|\Omega_k|^2} = \frac{1}{4} \left( 1 + \frac{|\Omega_k^+|}{|\Omega_k|} \right) \left( 1 - \frac{|\Omega_k^+|}{|\Omega_k|} \right) = \frac{1}{4} \left( 1 - \left( \frac{|\Omega_k^+|}{|\Omega_k|} \right)^2 \right).$$

Then Lemma 3.21 yields

$$\lim_{k \rightarrow \infty} \text{var}_k(f) = \frac{1}{4} \left( 1 - \frac{1}{(2\phi + 1)^2} \right).$$

Plugging in the value of  $\phi = (\sqrt{5} + 1)/2$  and simplifying completes the proof of the claim.  $\square$

**Theorem 3.24.** *The relaxation time and mixing time of the edge-flip Markov chain for dyadic tilings of size  $n$  are both at least  $\Omega(n^{2 \log \phi})$ , where  $\phi = \frac{\sqrt{5}+1}{2}$  is the golden ratio.*

*Proof.* We will derive an upper bound on the spectral gap  $\gamma_k$  using the distinguishing statistic  $f : \Omega_k \rightarrow \{0, 1\}$  given in Equation 3.3. We will apply this function to the characterization of the spectral gap in Proposition 2.2.

Because we have already shown that the variance of  $f$  is bounded away from 0 as  $k \rightarrow \infty$  (Lemma 3.23), it only remains to obtain an upper bound for  $\mathcal{E}(f)$ . Let  $\partial\Omega_k^+$  be the set of tilings in  $\Omega_k \setminus \Omega_k^+$  which can be obtained from a tiling in  $\Omega_k^+$  via one edge flip. Recall for two tilings  $x, y \in \Omega_k$ , we write  $x \sim_e y$  if  $x$  can be obtained from  $y$  by one edge flip. Hence,

$$\mathcal{E}(f) = \sum_{x \in \partial\Omega_k^+} \sum_{y \in \Omega_k^+ : y \sim_e x} \frac{1}{|\Omega_k|} \frac{1}{2n}.$$

Note that each tiling in  $\partial\Omega_k^+$  has a horizontal bisector and is not in  $\Omega_k^+$ . This means that it has exactly one edge flip that can bring it into  $\Omega_k^+$ , which is the flip that creates a vertical bisector. Then, we have

$$\mathcal{E}(f) = \frac{|\partial\Omega_k^+|}{2n \cdot |\Omega_k|}.$$

Now we need to describe the set  $\partial\Omega_k^+$ . It is a set of tilings with no vertical bisector, but with

one edge flip that creates a vertical bisector; see Figure 3.11. Note that the edge whose flip




Figure 3.11: A tiling in  $\partial\Omega_k^{\downarrow}$ , with the red edge being the flip that brings the tiling into  $\Omega_k^{\downarrow}$ .

creates a vertical bisector must be a horizontal edge of length 1 which flips to a vertical edge of length  $2/n$ . From now on we will refer to this edge as the *pivotal edge*.

In order to estimate the cardinality of  $\partial\Omega_k^{\downarrow}$ , we will describe a procedure to construct a tiling  $x \in \partial\Omega_k^{\downarrow}$ , observing the position of the pivotal edge. Note that  $x$  must have a horizontal bisector, which splits  $[0, 1]^2$  into its top and bottom halves. Assume that the pivotal edge is in the top half of  $x$ . This implies that the bottom half of  $x$  must itself contain a vertical bisector since the pivotal edge must be the only edge that forbids a vertical bisector to exist, see Figure 3.12(a). The two quadrants in the bottom half are simply any tilings of  $\Omega_{k-2}$ . Note also that the top half of  $x$  must contain a horizontal bisector, otherwise  $x \notin \partial\Omega_k^{\downarrow}$ , see Figure 3.12(b). Then we iterate the above construction: among the two halves of the top half, one must contain the pivotal edge, say the bottom one, while the other contains a vertical bisector, each side of which being completed with a tiling from  $\Omega_{k-3}$ , which gives the configuration in Figure 3.12(c). Continuing this for  $k - 2$  steps concludes the construction.

To estimate the cardinality of  $\partial\Omega_k^{\downarrow}$ , note that in each step of the construction we have two choices for where the pivotal edge is: either in the top half or the bottom half of the




Figure 3.12: The construction of a tiling in  $\partial\Omega_k^|$ . The grey areas represent the part that contains the pivotal edge.

corresponding region. Therefore, the number of tilings in  $\partial\Omega_k^|$  is

$$|\partial\Omega_k^| = \prod_{i=2}^k (2|\Omega_{k-i}|^2) = 2^{k-1} \prod_{i=0}^{k-2} |\Omega_i|^2 = \frac{n}{2} \prod_{i=0}^{k-2} |\Omega_i|^2.$$

Hence,

$$\mathcal{E}(f) = \frac{1}{4|\Omega_k|} \prod_{i=0}^{k-2} |\Omega_i|^2 \leq \frac{1}{4} \phi^{-2k+2}$$

where the last step follows from Lemma 3.22. Therefore, because  $\text{var}_k f$  converges to a positive constant (Lemma 3.23), there exists a constant  $c > 0$  such that

$$\gamma_k \leq \frac{\mathcal{E}(f)}{\text{var}_k f} \leq c\phi^{-2k}.$$

This implies that the relaxation time and mixing time satisfy

$$t_{\text{rel}}, t_{\text{mix}} \geq \frac{1}{c} \phi^{2k} = \frac{1}{c} \phi^{2 \log n} = \frac{1}{c} n^{2 \log \phi} = \Omega(n^{2 \log \phi}) = \Omega(n^{1.38}).$$

This complete the proof of the theorem. □

### 3.6 Weighted Dyadic Tilings and Statistical Physics

It is a general principle in statistical physics that in systems with some bias parameter (*temperature*) that induces different phases, the mixing time of natural heat-bath dynam-

ics should be as fast as possible (the diameter of the state space) at high temperature, a larger polynomial at the critical temperature, and exponential at low temperature. However, there are very few instances for which this behavior has been rigorously confirmed. Exceptions are the Ising model on complete graphs [77, 48], regular trees [47], and the two-dimensional lattice [78], and the Potts model on the complete graph [33] and the two-dimensional lattice [60], all of which required significant effort to analyze. The edge-flip Markov chain for dyadic tilings is an example of heat-bath dynamics, and our parameter  $\lambda$  can be viewed as a function of inverse temperature. The work in Sections 3.3 and 3.4 confirms exponential mixing at low temperature ( $\lambda > 1$ ) and polynomial mixing at high temperature ( $\lambda < 1$ ). Our work in Section 3.5 shows that the mixing time at the critical point ( $\lambda = 1$ ) is indeed polynomial but strictly larger than the diameter of the state space (which is  $n \log(n)/2$ ), providing further evidence for the general statistical physics principle above.

Another aspect of this general statistical physics principle is that the behavior at the critical point differs from the behavior near the critical point. We can demonstrate this is the case for  $\mathcal{M}_k^{edge}$  by appealing to relaxation times. We showed the relaxation time when  $\lambda < 1$  is  $O(n)$  and when  $\lambda = 1$  it is  $\Omega(n^{1.38})$  and  $O(n^{4.09})$  (Theorems 3.15, 3.24, and 3.20). It is a simple consequence of Theorem 2.5 and the relation between mixing and relaxation times (Proposition 2.1) that the relaxation time is at least exponential in  $n^2$  when  $\lambda > 1$ . We expect this this separation between the behavior at the critical point and behavior nearby also exists for mixing times, but our current results do not provably guarantee this.

### 3.7 Rectangular Dissections

In [24], we also considered rectangular dissections, tilings of the  $n \times n$  square with  $n$  rectangles, each of area  $n$ , whose corners lie in  $\mathbb{Z}^2$ . When scaled down by a factor of  $n$ , rectangular dissections are a generalization of dyadic tilings, with the dyadic constraint removed. We expected to see the same behavior for the edge-flip Markov chain for rectan-

gular dissections as we do for dyadic tilings, but the picture turned out to be significantly more complicated.

To begin, even showing that edge-flip moves connect the state space of all rectangular dissections required significant effort; this proof occupies about 10 pages in [24]. In fact, for a given rectangular dissection, it is not even obvious that a single valid edge-flip move exists. In contrast to our results for dyadic tilings, we proved this chain mixes in exponential time both when  $\lambda > 1$  and when  $\lambda < 1$ , though the reasons for slow mixing in the two regimes are different. It remains an open problem to bound the mixing time of this chain when  $\lambda = 1$ ; simulations seem to suggest it is possible this chain mixes in polynomial time at this isolated point.

All of this provides evidence that the dyadic constraint we place on our rectangles plays a critical role in determining the behavior of  $\mathcal{M}_k^{edge}$ . We depended on properties of dyadic tilings in our proofs when  $\lambda \leq 1$ , and this complementary work on rectangular dissections supports the notion that this is necessary, at least when  $\lambda < 1$ .

## CHAPTER 4

### FREE BOUNDARY PLANAR LATTICE PROBLEMS

Many statistical physics models are defined on an infinite lattice by taking appropriate limits of finite lattice regions, where a key consideration is how the boundaries are defined. For several models on finite planar lattices, such as 3-colorings and lozenge tilings, efficient sampling algorithms are known for regions with *fixed boundary conditions*, where the colors or tiles around the boundary are pre-specified [79], but much less is known about how to sample when these regions have *free boundaries*, where we want to include all configurations one could see within a finite window. It is widely believed that these chains are fast mixing on finite regions for any boundary conditions – that is, the emergent behavior has no dependence on the boundary conditions and there is only one phase – but previous proof techniques have been insufficient to rigorously verify this.

In this chapter, we introduce the new method of *randomized extensions* [25] and use it to extend the class of regions and boundary conditions for which efficient sampling is provably possible. At a high level, we relate sampling problems on regions with free boundaries or *mixed boundaries*, where some boundary colors might be fixed but others are allowed to vary, to a constant number of sampling problems on larger regions with fixed boundaries (for which efficient sampling processes are already known [79]). We use this principled approach to sample 3-colorings of a larger class of regions of  $\mathbb{Z}^2$  than was previously possible (Section 4.2) and show that the same approach can also give new results for sampling lozenge tilings of regions of the triangular lattice (Section 4.3). Sampling for these new types of regions is significant because it allows us to establish self-reducibility and gain a new ability to approximately count. We can then use these approximate counting algorithms as subroutines to enable sampling from an even larger class of regions. This makes progress towards the goal of showing that the emergent behavior (specifically, the mix-



ing time) of the sampling processes for planar lattice problems that we consider has no dependence on its boundary conditions.

At a high level, our new *randomized extensions* approach takes a free or mixed boundary planar region  $R$  and samples using a Markov chain that at each iteration:

- extends  $R$  three units in every direction to get a larger (fixed boundary, locally planar) region  $R'$ ;
- sets the configuration in  $R' \setminus R$  by choosing uniformly at random among a constant number of choices;
- makes a move of the (well-studied) related fixed boundary Markov chain in  $R'$ ; and
- restricts the resulting configuration back to  $R$ .

We make this principled approach concrete by applying it first to 3-colorings of finite subsets of  $\mathbb{Z}^2$  and then to lozenge tilings of finite planar regions. That this same approach yields new sampling results for different structures speaks to its generality.

## 4.1 Background on Grid 3-colorings

A *grid 3-coloring* is an assignment of one of three colors to each vertex in a subset of  $\mathbb{Z}^2$  such that no adjacent vertices are given the same color. We are interested in questions of sampling grid 3-colorings of finite, simply connected regions subject to various boundary constraints. Throughout we will label the three colors as  $\{0, 1, 2\}$ .

### 4.1.1 Related Work: Sampling $k$ -colorings

Colorings of the Cartesian lattice are of interest in statistical physics to study phase transitions of the  $k$ -state anti-ferromagnetic Potts model, a basic model of anti-ferromagnetism. In this model, vertices are assigned “spins” from  $\{1, \dots, k\}$ ; in the anti-ferromagnetic case,

neighbors have a preference for unequal spins. At zero temperature this preference becomes required, mapping the valid configurations precisely to the set of proper  $k$ -colorings, each occurring with equal probability.

Sampling proper  $k$ -colorings uniformly, from grids and from other graphs, has been the focus of much research. Given a graph  $G = (V, E)$  and an integer  $k$ , a  $k$ -coloring is an assignment of colors  $[k] = \{1, \dots, k\}$  to the vertices so that all pairs of neighboring vertices have distinct colors. A natural local Markov chain known as Glauber dynamics starts with any valid coloring, chooses  $(v, c) \in V \times [k]$  uniformly, and recolors  $v$  with color  $c$  if this yields a valid coloring.

Glauber dynamics have been extensively studied, primarily in the case when  $k$ , the number of colors/spins, is large compared to  $\Delta$ , the maximum degree of  $G$ . If  $k \geq \Delta + 2$  the chain is known to connect the state space. Vigoda [114] showed that the chain mixes in polynomial time when  $k \geq 11\Delta/6$ . For graphs with large girth and large  $\Delta$ , this degree constraint can be reduced [65, 75, 91]. Buble et al. [17] showed that the chain is rapidly mixing when  $k \geq 5$  whenever  $\Delta = 3$  or when  $k \geq 7$ ,  $\Delta = 4$  and  $G$  is triangle free, notable because this includes the Cartesian lattice  $\mathbb{Z}^2$ . See [59] for a survey on efficiently sampling  $k$ -colorings.

#### 4.1.2 Sampling Grid 3-Colorings

For a small number of colors, Glauber dynamics has proven more challenging to analyze. Despite this, some results are known for 3-colorings on  $\mathbb{Z}^2$ , which map bijectively to Eulerian orientations of the grid graph, well-known in the statistical physics community as the “ice model.” This structure allows more analysis, and Luby et al. [79] showed that a related Markov chain based on “tower moves” that can at once update a linear collection of sites is rapidly mixing on any simply connected region of  $\mathbb{Z}^2$ , provided the configuration on the boundary is fixed in advance (known as *fixed boundary conditions*). Randall and Tetali [97] subsequently showed polynomial convergence of Glauber dynamics itself could

be inferred from the comparison method of Diaconis and Saloff-Coste [43]. Though the results of [79] and [97] are presented in the context of Eulerian orientations, fast mixing of Glauber dynamics for 3-colorings of finite simply connected subsets of  $\mathbb{Z}^2$  with fixed boundary colors is an immediate corollary.

There remain basic open questions about the convergence of Glauber dynamics for sampling 3-colorings in  $\mathbb{Z}^2$ , particularly in the context of *free boundary conditions*. This question is computationally interesting because for some related models the convergence rates are known to depend significantly on the types of boundaries (see, e.g., [11, 84, 85]). Goldberg, Martin and Paterson [61] extended the Markov chain studied by Luby et al. to the case of free boundary conditions on rectangular subregions of  $\mathbb{Z}^2$ , but their argument does not seem to extend to other simply connected regions. While rectangular regions are of the most significance in physics, the restriction to this class of graphs precludes, for instance, “L-shaped” regions that are necessary for the self-reducibility that allows us to approximately count using a well-known reduction between sampling and approximate counting due to Jerrum, Valiant and Vazirani [71]. Moreover, the proof of Goldberg et al. [61] identifies a set of weights for towers moves (to be defined in Section 4.1.4) near or intersecting the boundary, and shows that these weights allow the coupling proof of Luby et al. [79] to extend to rectangular regions with free boundary conditions. There is no explanation for why their weights work or why a similar approach apparently fails for more general lattice regions.

#### 4.1.3 Problem Statement and Assumptions

Formally, we assume we are given as input a simply connected finite subset  $R$  of  $\mathbb{Z}^2$  with some (possibly zero) colors adjacent to the boundary fixed such that  $R$  admits a 3-coloring. We want to generate, uniformly at random, one such 3-coloring of  $R$  that respects the values of these fixed cells. For simplicity, we assume  $R$  is a subset of the Cartesian grid and we are assigning colors to faces (rather than to vertices). For example, see Figure 4.1a for

2						
1						
2	0					
	1					
	2					

(a)

2	1	2	0	1	2	0
1	2	0	2	0	1	2
2	0	1	0	1	2	0
	1	0	1	2	0	1
	2	2	0	1	2	0

(b)

Figure 4.1: (a) A sample input: a mixed boundary region  $R$  for which we want to efficiently generate a uniformly random 3-coloring, and (b) one such proper 3-coloring of  $R$ .

an example of an input we will consider, and Figure 4.1b for a sample we might generate given this input. Throughout we will let  $|R|$  denote the number of cells in  $R$  without a fixed color; in Figure 4.1,  $|R| = 27$ . We assume all cells of  $R$  without fixed colors are connected.

By a *convex corner cell* of a grid region  $R$ , we mean a cell where two of its adjacent sides are part of  $\partial R$ , the boundary of  $R$ . By a *reflex cell*, we mean a cell of  $R$  where exactly one corner but no sides are in  $\partial R$ ; equivalently, a reflex cell has all horizontal and vertical neighbors in  $R$  but at least one diagonal neighbor not in  $R$ . In Figure 4.1a, the cells given color 2 are convex corner cells and the cell given color 0 is a reflex cell. While not necessary, we assume throughout that  $R$  has width at least two everywhere, e.g., there are no cells where two opposite sides are both included in different segments of  $\partial R$ . Because we will always assume that any reflex cells of  $R$  have fixed colors, any such channel of width 1 must end at either a fixed color cell or at a boundary of  $R$ . Sampling 3-colorings of a  $1 \times k$  region that has no influence on the coloring elsewhere is easy to do, so we eliminate this case to simplify our analysis. Specifically, by eliminating  $1 \times k$  regions from consideration, we assume that all cells of  $R$  are interior cells, adjacent to one boundary segment, adjacent to two boundaries at a convex corner, or are reflex cells. Throughout, by a *grid region*  $R$  we will mean a simply connected finite region of  $\mathbb{Z}^2$  that satisfies this condition.

#### 4.1.4 Glauber Dynamics and Tower Moves

Glauber dynamics is a type of local update Markov chain that has been widely studied, particularly in statistical physics. For 3-colorings, Glauber dynamics pick a random cell  $v \in R$  (that does not have a fixed color) and a random color  $c \in \{0, 1, 2\}$ , and recolors  $v$  with  $c$  if doing so yields a valid 3-coloring. We call such an update of one cell to a new color a *flip move*. It has proven difficult to construct direct proofs bounding the convergence times of Glauber dynamics on various planar lattice structures such as 3-colorings and lozenge tilings.

Luby et al. [79] instead introduced “tower-based Markov chains” that sample from these structures efficiently by updating multiple locations at once. Fast mixing of this chain implies fast mixing of Glauber dynamics [97]. We take the same approach, bounding the mixing time of a tower Markov chain which, by the comparison method [43], gives bounds on Glauber dynamics. We begin by defining tower moves.

**Definition 4.1.** *A tower in a 3-coloring of a region  $R$  is a set of contiguous cells in a row or column that can be labelled in order (right to left, left to right, top to bottom, or bottom to top) as  $v_1, v_2, \dots, v_h$  for some  $h \geq 2$  such that:*

- *For all  $i = 1, 2, \dots, h - 1$ ,  $v_i$ 's neighbors except for  $v_{i+1}$  are the same color*
- *$v_h$  has all neighbors of the same color.*

*We say that  $v_1$  is the start of the tower,  $v_h$  is the end of the tower, and the height of the tower is  $h$ .*

**Definition 4.2.** *A tower move recolors a tower such that for  $i = 1, 2, \dots, h - 1$ ,  $v_i$  receives the color of  $v_{i+1}$ , and  $v_h$  receives the color distinct from its original color and its neighbors. The start color of a tower move is the color of  $v_2$ , the color  $v_1$  is given after the tower move.*

A tower move can be thought of as attempt to change the tower's start cell to its start color  $c$  that, at the same time, forces other cells to also change their colors so that this




Figure 4.2: A 3-coloring of a rectangle  $R$  and three tower moves, each of height 3. The top right tower abuts the boundary, the bottom tower is adjacent to the boundary, and the top left tower does both. The arrow drawn in each tower goes from the start cell to the end cell.

recoloring of the start cell is valid. See Figure 4.2 for examples of different towers of height 3. Note that if a tower includes a fixed boundary cell, a tower move is not possible because fixed color boundary cells cannot be given a different color.

We will denote the boundary of region  $R$  by  $\partial R$ . A tower *abutting* a segment  $\ell$  of  $\partial R$  stretches perpendicular to  $\ell$  such that its start or end cell is adjacent to  $\ell$ . A tower *adjacent* to a segment  $\ell$  of  $\partial R$  stretches parallel to  $\ell$  and has all of its cells next to  $\ell$ . Figure 4.2 shows towers that are abutting a boundary (top right), adjacent to a boundary (bottom), and both (top left). We consider 3-colorings of a variety of non-convex regions, but the conditions we impose ensure these remain the only types of towers we need to consider.

#### 4.1.5 Height Functions

It is well-known [61, 79] that grid 3-colorings can be mapped to height functions. A *height function* for a 3-coloring of a grid region  $R$  is an assignment of an integer (a *height*) to each cell in  $R$  such that two adjacent cells in  $R$  differ in height by exactly one and the height of a cell (modulo 3) is its color.

More formally, let  $\sigma$  be a 3-coloring of  $R$ , where for each cell  $b \in R$ ,  $\sigma(b) \in \{0, 1, 2\}$ . We can define a height function  $h_\sigma : R \rightarrow \mathbb{Z}$  by picking any cell  $a \in R$  that has color  $\sigma(a) = i$  and setting  $h_\sigma(a) = i$ , and (uniquely) completing the heights for all remaining

2	1	2	0	1	2	0
0	2	0	2	0	1	2
1	0	1	0	1	2	0
0	2	0	1	2	0	1
1	0	2	0	1	2	0

(a)

-1	-2	-1	0	1	2	3
0	-1	0	-1	0	1	2
1	0	1	0	1	2	3
0	-1	0	1	2	3	4
1	0	-1	0	1	2	3

(b)

Figure 4.3: (a) a valid 3-coloring of a  $7 \times 5$  rectangle, and (b) the height function obtained from it when fixing the height of the lowest leftmost cell to be its color.

cells subject to the rules

$$h_\sigma(x, y) - h_\sigma(x - 1, y) = \begin{cases} 1 & \sigma(x, y) - \sigma(x - 1, y) \equiv 1 \pmod{3} \\ -1 & \sigma(x, y) - \sigma(x - 1, y) \equiv -1 \pmod{3} \end{cases}$$

$$h_\sigma(x, y) - h_\sigma(x, y - 1) = \begin{cases} 1 & \sigma(x, y) - \sigma(x, y - 1) \equiv 1 \pmod{3} \\ -1 & \sigma(x, y) - \sigma(x, y - 1) \equiv -1 \pmod{3} \end{cases}$$

There may be many height functions depending on the chosen a starting cell  $a$ , but for any two  $h_\sigma$  and  $h'_\sigma$  there is a constant  $k$  such that  $h_\sigma(b) - h'_\sigma(b) = 3k$  for all  $b \in R$ . A 3-coloring and the height function obtained from it when fixing the height of the lowest leftmost cell to be its color are shown in Figure 4.3.

The most important use of height functions for 3-colorings has been in understanding how to move from one 3-coloring of a region to another using flip moves. We use height functions in Section 4.2.6 to characterize when our Markov chain  $\mathcal{M}_C$  for sampling 3-colorings is irreducible, and in Section 4.2.7 in our path coupling proof of fast mixing for  $\mathcal{M}_C$  to give an upper bound on the number of flip moves required to move from one coloring to another. We now present some preliminaries about height functions that will be useful in these two contexts.

**Definition 4.3.** *Given a grid region  $R$  with mixed boundary conditions, we say that the boundary conditions are height consistent if the heights of all fixed cells in any valid 3-*




Figure 4.4: (a) Three regions with height consistent mixed boundary conditions, where gray cells have fixed colors. (b) A region whose mixed boundary conditions are not height consistent; in the two colorings shown, the fixed cells have different heights.

*coloring can be uniquely determined from the height of any one such cell.*

For example, when the fixed color cells are contiguous along the boundary of  $R$ , these mixed boundary conditions are necessarily height consistent. If there are zero or one fixed color cells on the boundary of  $R$ , vacuously there is height consistency. Figure 4.4 gives more examples of height consistent and non-height consistent boundaries. Note that mixed boundary conditions must be height consistent if Glauber dynamics connects the state space, because Glauber dynamics cannot change the heights associated to any fixed cells; we'll later prove these two conditions are in fact equivalent.

We now use height functions to give an upper bound on the number of flips needed to move between any two 3-colorings of a mixed boundary region  $R$ ; these upper bounds will be used in our path coupling argument in Section 4.2.7.

**Lemma 4.4.** *Let  $R$  be a mixed boundary region with at least one fixed color cell, and let  $\overline{\Omega}$  be the set of valid 3-colorings of  $R$  reachable by flips from some initial 3-coloring  $\sigma$ . The number of flips required to move between any two 3-colorings of  $R$  is at most  $|R|^2$ , where  $|R|$  is the number of cells in  $R$  without fixed colors.*

*Proof.* Pick one fixed boundary cell  $a$  of  $R$ , and suppose its color is  $c$ . For each 3-coloring  $\sigma$  in  $\overline{\Omega}$ , let  $h_\sigma$  be the height function for  $\sigma$  that sets  $h_\sigma(a) = c$ . For any two height functions  $h_\sigma$  and  $h_\tau$  in  $\overline{\Omega}$ , define the *area* between them to be  $\sum_{b \in R} |h_\sigma(b) - h_\tau(b)|$ . If  $b$  is a fixed color cell of  $R$  then because  $\overline{\Omega}$  is connected by flips it must be that  $h_\sigma(b) = h_\tau(b)$  because flip moves do not change heights of fixed boundary cells, so only free cells contribute to



the area between two height functions. We note that free cells  $b$  at distance  $i$  from  $a$  have  $|h_\sigma(b) - h_\tau(b)| \leq 2i$ . Each of the  $|R|$  free cells of  $R$  is at distance at most  $|R|$  from  $a$ , so we have that the area between  $h_\sigma$  and  $h_\tau$  is at most  $2|R|^2$ . It is a well-known fact that for any two 3-colorings that are connected by flips, as all tilings in  $\overline{\Omega}$  must be, there exists at least one flip move that decreases the area between their associated height functions. Each such flip move in fact decreases this area by two. This implies the number of flip moves needed to move between any two tilings of  $\overline{\Omega}$  is at most  $2|R|^2/2 = |R|^2$ .  $\square$

We note effort could be made to improve the bound in this lemma to about  $|R|^2/2$ , but this has no asymptotic effect on our results.

The above lemma requires the presence of at least one fixed boundary cell. Because we will assume all reflex cells of  $R$  have fixed colors, the only case not yet addressed is free boundary regions with no reflex cells, which are necessarily rectangles. Free boundary rectangles are exactly the case considered by [61], and we briefly present their similar result.

**Lemma 4.5** ([61]). *Let  $R$  be a free boundary rectangle. Then the number of flips needed to get from any 3-coloring of  $R$  to any other 3-coloring of  $R$  is at most  $|R|^2$ .*

*Proof.* The authors of [61] prove that for  $n \times m$  rectangles with  $m \geq n$ , the number of flips required to get from any 3-coloring to any other 3-coloring is at most  $2nm^2$ . Since  $|R| = mn$ , and by assumption (Section 4.1.3) grid region  $R$  has length and height both at least 2, then  $|R|^2$  is an upper bound on this number.  $\square$

## 4.2 A Markov Chain for Grid 3-colorings using Randomized Extensions

In this section we present and analyze a tower-based Markov chain for uniformly sampling 3-colorings of finite simply connected subsets of the Cartesian grid with mixtures of free and fixed boundaries; using standard techniques [43], one can extend rapid mixing results for this tower chain to Glauber dynamics. Our tower Markov chains uses randomized

extensions to generate, from a mixed boundary region  $R$ , a larger fixed boundary region  $R'$ , does a move of the tower Markov chain from [79] for fixed boundary 3-colorings on  $R'$ , restricts the resulting configuration back to  $R$ , and repeats; we formally define this chain in Section 4.2.2.

We will show for 3-colorings on any simply connected subset of the Cartesian grid with height consistent mixed boundary conditions that this tower-based Markov chain using randomized extensions converges to the uniform distribution in polynomial time, provided all reflex cells have fixed colors. For example, if first we flip a coin to determine the color of the single reflex cell in any “L-shaped” region of  $\mathbb{Z}^2$ , and then run Glauber dynamics on the remaining cells, our proof shows this chain will converge quickly; this is the first result to efficiently sample uniformly from free boundary 3-colorings of any non-convex region.

Another important class for which our work shows fast convergence is L-shaped regions with mixed boundaries, where the fixed part is contiguous and includes the unique reflex cell. This class of regions is precisely what is required to establish self-reducibility of the sampling problem. We can thus use sampling to, for the first time, efficiently approximately count the number of 3-colorings of any rectangular region with free boundary conditions utilizing the technique of Jerrum, Valiant, and Vazirani [71]. In fact, these approximate counting results hold for any height-consistent mixed boundary regions where all reflex cells have fixed colors.

#### 4.2.1 Techniques

Our proofs follow a similar strategy as Goldberg et al. [61]; however, rather than checking whether there are probabilities for tower moves near free boundaries that enable a path coupling argument, we *derive* these probabilities using *randomized extensions*. At each iteration, we embed the region of interest with free or mixed boundaries into a larger region with fixed boundaries by extending every boundary segment by three cells. The colors within this extension are randomized so that, at each iteration, for any valid move near a

free boundary there is at least one coloring of the extension in which this move is valid for the extended region. While it is possible to guarantee that this occurs by considering all (exponentially many) colorings of the extended region, perhaps surprisingly we show that randomly choosing from merely a *constant number* of colorings of the extension suffices. We then use path coupling arguments regarding tower Markov chains for fixed boundary regions [79] to bound the mixing time of our Markov chain on free and mixed boundary regions. A comparison argument allows us to infer that Glauber dynamics also mixes in polynomial time. Because handling extensions near reflex corners is more challenging, we require that the color at any reflex corner is fixed before sampling occurs.

An interesting feature of this new approach is that we don't need to know a priori the probabilities of moves within  $R$  occurring near free boundaries. Instead, we derive these probabilities as a result of the randomized extensions we generate. For 3-colorings on rectangular regions of  $\mathbb{Z}^2$  with free boundaries, the probabilities we derive exactly coincide with those used by [61]. While the authors of [61] simply stated, without justification, the probabilities of moves near free boundaries and showed they enabled fast mixing, our principled approach of deriving these probabilities allows us to generalize our arguments to more regions and mixed boundary conditions, as well as to lozenge tilings.

#### 4.2.2 A Sampling Algorithm using Randomized Extensions

We now state our Markov chain algorithm  $\mathcal{M}_C$  for sampling from 3-colorings of a simply connected Cartesian grid region  $R$  that has some (possibly zero) boundary cells with fixed colors. We assume the color at any reflex corner of  $R$  is fixed. This Markov chain is stated as Algorithm 4, where Step 1 – how random extensions are generated – is purposefully vague; in the next sections we will explore this step in more detail.

Throughout we let  $P_C$  denote the transition matrix of this Markov chain. That is, by  $P_C(\sigma, \tau)$  we mean the probability, if at 3-coloring  $\sigma$ , of going to  $\tau$  in one step. We note that the probability of proposing a given move in  $R'$  is  $1/|R'| \cdot 1/3$ , where  $|R'|$  is the number of

---

**Algorithm 4** Markov Chain  $\mathcal{M}_C$  on grid region  $R$  with mixed boundary conditions

---

Starting at any initial 3-coloring  $\chi_0$  of  $R$ , repeat:

- 1: Choose, uniformly at random, an extension of coloring  $\chi_i$  of  $R$  to a coloring  $\chi'_i$  of larger region  $R'$ .
  - 2: Choose, uniformly at random, a cell  $a$  of  $R'$  that does not have a fixed color; a color  $c \in \{0, 1, 2\}$ ; and a probability  $p \in (0, 1)$ .
  - 3: **if**  $\chi'_i(a) \neq c$  and no neighbors of  $a$  have color  $c$  in  $\chi'_i$  **then**
  - 4:     recolor  $a$  with color  $c$ .
  - 5: **else if**  $a$  starts a tower of color  $c$  that doesn't include any fixed cells and whose intersection with  $R$  has size  $h \geq 1$  **then**
  - 6:     make this tower move if  $p < 1/2h$ .
  - 7: Let  $\chi_{i+1}$  be the resulting coloring of  $R$ .
- 

cells in  $R'$  without fixed colors. We now explore Step 1 of this algorithm, how to generate and color the random extension  $R'$ , in more detail.

#### 4.2.3 Extending One Boundary Segment

Consider any straight line segment  $\ell$  of the boundary  $\partial R$  of  $R$ . Some colors along  $\ell$  may be fixed, and some may not be. In this section we explore how to randomly extend the boundary of  $R$  past  $\ell$  to form  $R'$ ; subsequent sections will deal with extensions near where two boundary segments of  $R$  meet at either a convex or reflex corner.

Throughout, fixed color cells extend to fixed color cells and free color cells extend to free color cells. It is superfluous to extend fixed color cells, except when they are adjacent to free color cells, because no moves can ever occur in fixed color regions. For simplicity we describe extensions for entire segments of  $\partial R$ , even though it suffices to only extend free color cells and any fixed color cells adjacent to them.

Without loss of generality, we suppose  $\ell$  is vertical and the interior of  $R$  lies to its left. To obtain  $R'$  near  $\ell$ , we extend  $R$  to the right three units (See Figure 4.5). For now we assume this extension does not intersect any other portions of  $R$ ; exceptions to this will be discussed later. Given a coloring  $\chi$  of the cells of  $R$  adjacent to  $\ell$ , we extend it in one of two ways, each with equal probability, to a coloring  $\chi'$  of the extended portion of  $R'$  beyond  $\ell$  as follows. Within  $R$ ,  $\chi' = \chi$ . Along the right side of  $\ell$ , the colors in the additional columns

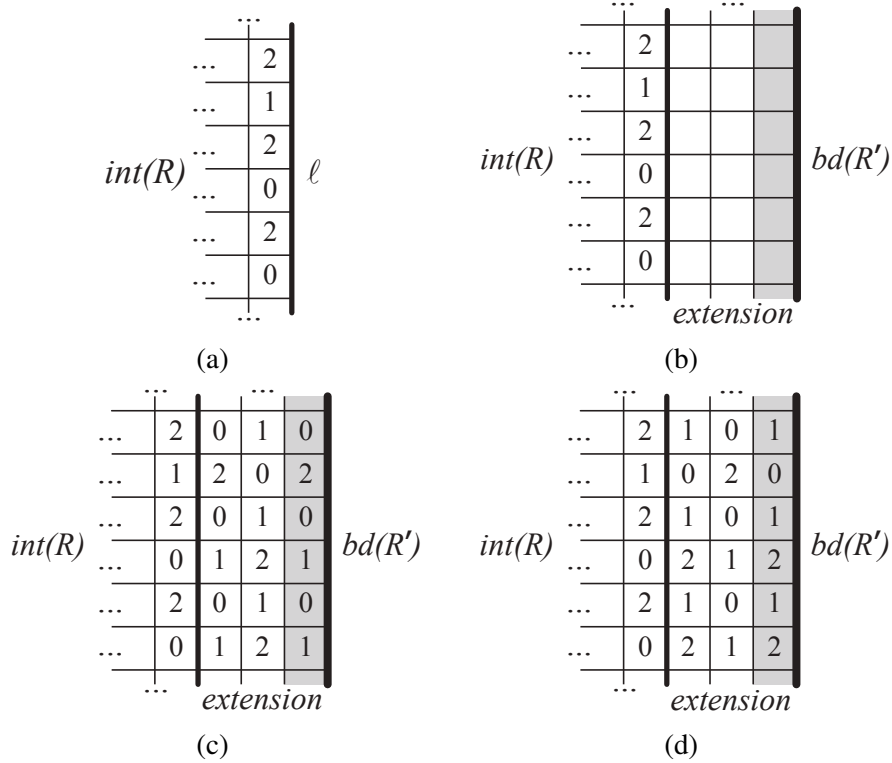


Figure 4.5: (a) A portion  $\ell$  of the boundary of a region  $R$  lined by free boundary cells (with no fixed colors). (b) To obtain  $R'$  we extend  $R$  three units past  $\ell$ . We will fix the colors adjacent to the boundary of this new region  $R'$ . (c)  $R$ 's up-up-down (UUD) coloring of the extended region  $R'$ . (d)  $R$ 's down-down-up (DDU) coloring of the extended region  $R'$ . Throughout gray cells have fixed color.

will each be a copy of the values of  $\chi$  in the column  $C$  of  $R$  adjacent to  $\ell$ , denoted  $\chi(C)$ , with the values incremented or decremented by 1 (mod 3). With probability 1/2, the colors of the first column to the right of  $\ell$  are  $\chi(C) + 1(\text{mod } 3)$ , the colors of the second column are  $\chi(C) + 2(\text{mod } 3)$ , and the colors of the third column are  $\chi(C) + 1(\text{mod } 3)$ ; see Figure 4.5c. The colors in these columns can be seen as  $\chi(C)$  incremented, then incremented again, and then decremented; we will refer to such a configuration as “up-up-down,” or UUD for short. With the remaining probability 1/2, the columns right of  $C$  will have a “down-down-up” (DDU) configuration, consisting of, left to right,  $\chi(C) - 1(\text{mod } 3)$ ,  $\chi(C) - 2(\text{mod } 3)$ , and  $\chi(C) - 1(\text{mod } 3)$ ; see Figure 4.5d. We then treat the colors along the boundary of this new region  $R'$  as fixed.

We now examine moves that occur in  $R'$  near a boundary segment, and how probabili-

ties of moves in  $R'$  translate into probabilities of moves when looking only within  $R$ . We focus on moves where each cell whose color is updated is adjacent to at most one boundary segment, as moves near corners where boundary segments meet will be considered in the next two sections.

**Lemma 4.6.** *Let  $\sigma$  and  $\tau$  be 3-colorings of a free or mixed boundary region  $R$  where any reflex cells have fixed colors. Suppose  $\tau$  can be obtained from  $\sigma$  in one iteration of  $\mathcal{M}_C$  via a move that changes cell(s) that touch most one segment of  $\partial R$ . Then,  $P_C(\sigma, \tau) = P_C(\tau, \sigma)$ .*

*Proof.* We define  $q := 1/(3|R'|)$  to be the probability of proposing a move by picking a location in  $R'$  (without a fixed color) and a desired update color. We briefly note that for every extension of  $R$ , the number  $|R'|$  of cells in  $R'$  without a fixed color is the same - all free boundaries are always extended three units, and the only differences between different random extensions are the colors used.

To begin, if all cells that change from  $\sigma$  to  $\tau$  are in the interior of  $R$ , then the extensions chosen play no role in the move and moves occur with exactly the probability they do in the original fixed boundary tower Markov chain of [79]. That is, if  $\sigma$  and  $\tau$  differ by a single interior flip, then

$$P_C(\sigma, \tau) = P_C(\tau, \sigma) = \frac{1}{3|R'|} = q$$

If  $\sigma$  and  $\tau$  differ by a single interior tower of height  $h$ , then

$$P_C(\sigma, \tau) = P_C(\tau, \sigma) = \frac{1}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{2h}$$

For moves from  $\sigma$  to  $\tau$  that satisfy the hypotheses of the theorem but are not in the interior of  $R$ , because we assume reflex cells have fixed colors and cannot be included in any move then there are only four cases to consider: flips changing the color of a cell next to one boundary segment  $\ell$ ; towers abutting one boundary segment  $\ell$ ; towers abutting two boundary segments  $\ell_1$  and  $\ell_2$ , one at each of its start and end cells; or towers *adjacent* to one boundary segment  $\ell$ . Avoiding towers with some intermediate number of cells adjacent

to a free boundary segment  $\ell$  is a main reason we require all reflex cells to have fixed colors throughout this section.

We first consider towers abutting one boundary as a detailed illustrative case, and then continue to analyze the remaining three cases.

*Case 1: Tower abutting one free boundary, height  $h \geq 2$ .* Suppose  $\tau$  can be obtained from  $\sigma$  via a tower move of height  $h \geq 2$  that abuts  $\ell$ ; see Figure 4.6a which shows a sample configuration  $\sigma$  and a tower move of height  $h = 3$  that yields  $\tau$  (shown in Figure 4.6b). We will calculate the probability of such a tower move transforming  $\sigma$  into  $\tau$  in one iteration of  $\mathcal{M}_C$ , as well as the probability of the reverse tower move transforming  $\tau$  into  $\sigma$ . Without loss of generality, let the tower's leftmost cell  $(\bar{x}_1, \bar{y})$  be its start cell, where  $\sigma(\bar{x}_1, \bar{y}) = 0$  and  $\tau(\bar{x}_1, \bar{y}) = 1$ , and suppose the tower's colors increase (mod 3) from left to right (all other cases can be obtained by permuting the colors or permuting  $\sigma$  and  $\tau$  in all arguments below). Let  $(\bar{x}_2, \bar{y})$  denote the tower cell abutting  $\ell$ ,  $(\bar{x}_2 + 1, \bar{y})$  its right neighbor (in the random extension) and  $(\bar{x}_2 + 2, \bar{y})$  the cell two units to its right. Let  $c_2$  be the color of cell  $(\bar{x}_2, \bar{y})$ ; in Figure 4.6,  $c_2 = 2$ .

From coloring  $\sigma$ , the transitions of  $\mathcal{M}_C$  yielding coloring  $\tau$  are tower moves that begin at  $(\bar{x}_1, \bar{y})$  with start color  $c = 1$  and  $p < 1/2h$ ; in one coloring of the extension beyond  $\ell$  the tower ends at  $\ell$  and in the other it continues to  $(\bar{x}_2 + 2, \bar{y})$  (see the top row of Figure 4.6c). Due to the nature of our extensions, all tower moves beginning in the interior of  $R$  always terminate before the fixed boundary of  $R'$  and thus do not contain one of  $R'$ 's fixed boundary cells. Even though this tower has a different height in  $R'$  depending on what the coloring of the extension is, both towers' intersection with  $R$  is the same height,  $h$ . We see that

$$P_C(\sigma, \tau) = \frac{1}{|R'|} \cdot \frac{1}{3} \cdot \frac{1}{2h} = \frac{q}{2h}.$$

From  $\tau$ , there are two tower transitions yielding  $\sigma$ , each occurring when the extension is up-up-down (UUD) but with different starting points; when the extension is down-down-up (DDU), cell  $(\bar{x}_2, \bar{y})$  has two neighbors of each color distinct from its own and cannot be

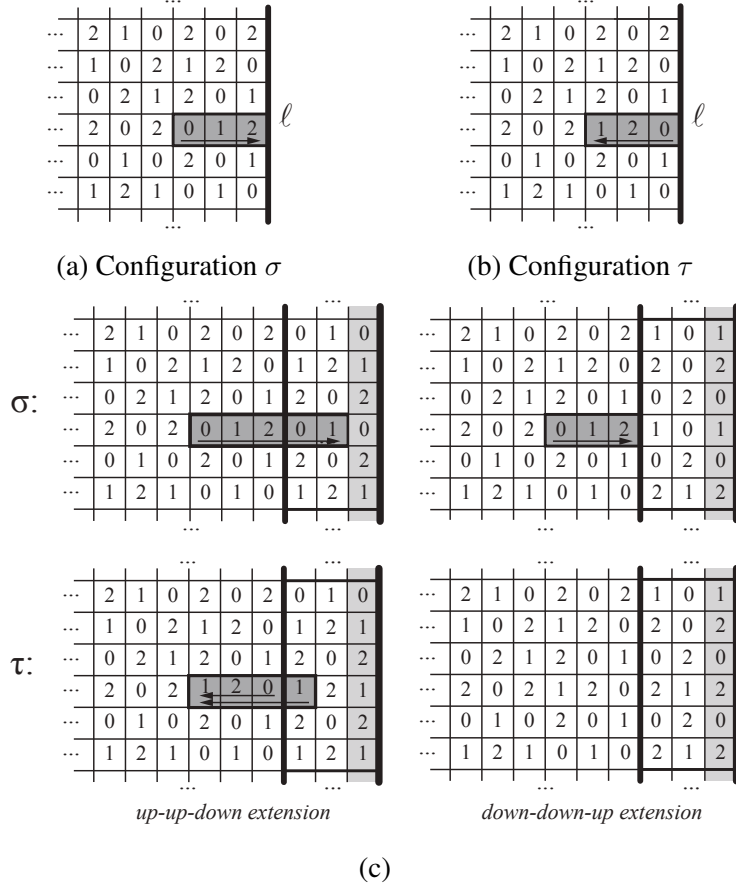


Figure 4.6: Colorings (a)  $\sigma$  and (b)  $\tau$  differ by a tower of height 3 (grey) abutting boundary segment  $\ell$ . (c) The indicated tower moves for the random colorings of the extended region  $R'$  transition between  $\sigma$  and  $\tau$ .

part of any tower (see the bottom row of Figure 4.6c). Specifically, the transitions of  $\mathcal{M}_C$  yielding  $\sigma$  correspond to choosing cell  $(\bar{x}_2, \bar{y})$  and color  $c_2$ , or choosing cell  $(\bar{x}_2 + 1, \bar{y})$  and color  $c_2 + 1 \pmod{3}$ . Both of these towers have the same height  $h$  within  $R$ , which is also the same height of the towers used to go from  $\sigma$  to  $\tau$ . To transform  $\tau$  into  $\sigma$ , we must choose the correct coloring of the extension beyond  $\ell$ , choose the correct starting point and color for one of the two towers, and have  $p < 1/2h$ :

$$P_C(\tau, \sigma) = \frac{1}{2} \cdot \frac{2}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{2h}.$$

We conclude that  $P_C(\sigma, \tau) = P_C(\tau, \sigma) = q/2h$  when  $\sigma$  and  $\tau$  differ by a tower abutting



one boundary segment  $\ell$ .

Case 2: Tower abutting two boundaries, height  $h \geq 2$ . If a tower abuts two boundaries, one at each end, then the configuration at each end of the tower is as described in Case 1. There are four possible extensions to consider, two for the boundary at each end of the tower. Suppose in  $\sigma$  the tower's start cell is adjacent to a boundary  $\ell_1$ . Then, just as in Case 1 for  $\tau$ , across the two extensions for  $\ell_1$  there are two possible choices of a start cell and a start color that would begin this tower move; it will be a valid tower move whose intersection with  $R$  is height  $h$  no matter the extension of  $\ell_2$ , the boundary adjacent to the tower's end cell; and the move will occur with probability  $1/2h$ . Thus, we get

$$P_C(\sigma, \tau) = \frac{1}{2} \cdot \frac{2}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{2h}.$$

Similarly, in  $\tau$  this tower's start cell will be adjacent to  $\ell_2$ ; across the two extension for  $\ell_2$  there are two possible start cell, start color pairs that can initiate the tower move; the tower move is valid for any extension of  $\ell_1$  and always has the same height within  $R$ ; and each move will occur with probability  $1/2h$ . That is,

$$P_C(\tau, \sigma) = \frac{1}{2} \cdot \frac{2}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{2h}.$$

We conclude that  $P_C(\sigma, \tau) = P_C(\tau, \sigma) = q/2h$  for any towers abutting two boundaries.

Case 3: Free boundary flip. Consider two 3-colorings  $\sigma$  and  $\tau$  of  $R$ , differing at a single location  $(\bar{x}, \bar{y})$  in  $R$  adjacent to  $\ell$ . Without loss of generality, suppose  $\sigma(\bar{x}, \bar{y}) = 1$  and  $\tau(\bar{x}, \bar{y}) = 2$ ; all other cases can be obtained by permuting the colors in all arguments below. Note this implies all three neighbors within  $R$  of location  $(\bar{x}, \bar{y})$  have color 0; see Figure 4.7.

We first consider moves that transform  $\sigma$  into  $\tau$ . When the extension for  $\sigma$  is up-up-down, cell  $(\bar{x}, \bar{y})$ 's neighboring location  $(\bar{x} + 1, \bar{y})$  outside of  $R$  is given color 2. Picking cell  $(m - 1, \bar{y})$  and color 2 then begins a right-going tower whose intersection with  $R$  is

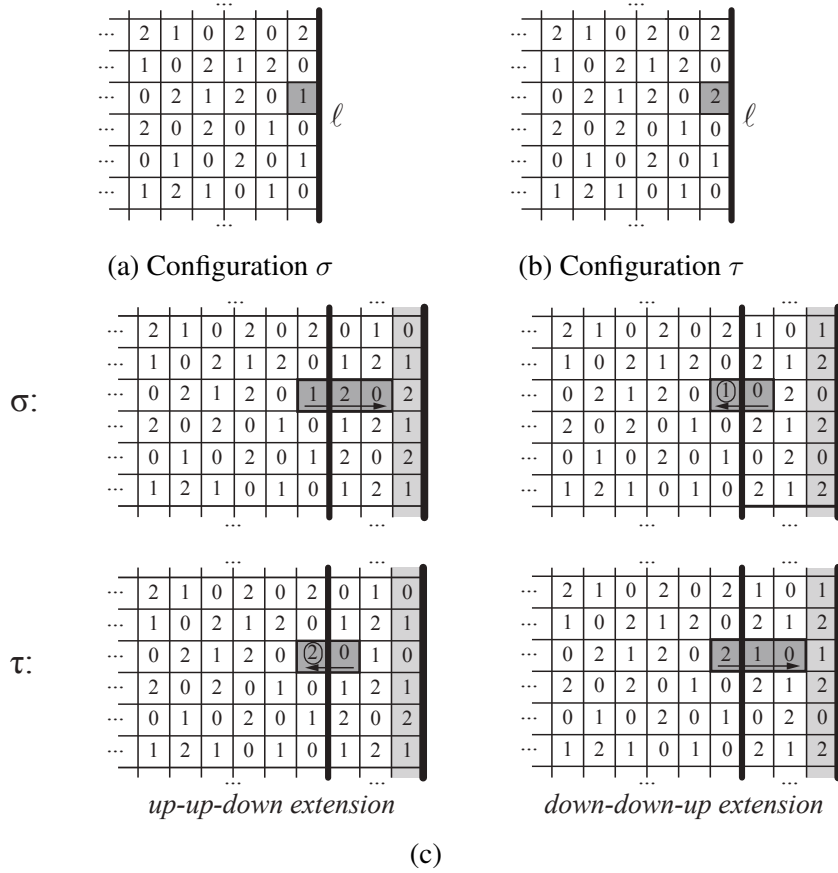


Figure 4.7: Colorings (a)  $\sigma$  and (b)  $\tau$  differ at a single cell (grey) next to a boundary segment  $\ell$ . (c) The indicated tower moves (dark gray, and indicated by arrows) and flip moves (indicated by circles) for the random colorings of the extended region  $R'$  transition between  $\sigma$  and  $\tau$ .

of height 1, so this tower move occurs and produces  $\tau$  with probability  $1/2$ . When the extension for  $\sigma$  is down-down-up, picking cell  $(\bar{x}, \bar{y})$  and color 2 initiates a valid flip move that results in  $\tau$ . In this DDU case there is also an additional tower move with start color 1 beginning in  $R'$  at  $(\bar{x} + 1, \bar{y})$  that ends at  $(\bar{x}, \bar{y})$  and changes  $\sigma$  into  $\tau$ ; because this tower's intersection with  $R$  is height 1, this tower move occurs with probability  $1/2$  if selected. We see that:

$$P_c(\sigma, \tau) = \frac{1}{2} \cdot \frac{1}{3|R'|} \cdot \left( \frac{1}{2} + 1 + \frac{1}{2} \right) = q.$$

Similarly, there exist three moves that can transform  $\tau$  to  $\sigma$ : In  $\tau$ 's DDU extension, a tower beginning at  $(\bar{x}, \bar{y})$  with start color 1 that happens with probability  $1/2$ ; in  $\tau$ 's UUD

extension, a tower beginning at  $(\bar{x} + 1, \bar{y})$  with start color 2 that occurs with probability  $1/2$ ; and again in  $\tau$ 's UUD extension, a flip move at  $(\bar{x}, \bar{y})$  to color 1. We see that

$$P_c(\tau, \sigma) = \frac{1}{2} \cdot \frac{1}{3|R'|} \cdot \left( \frac{1}{2} + \frac{1}{2} + 1 \right) = q.$$

We conclude that  $P_C(\sigma, \tau) = P_C(\tau, \sigma) = q$ .

Case 4: Tower adjacent to the free boundary, height  $h \geq 2$ . Lastly, consider two 3-colorings  $\sigma$  and  $\tau$  of  $R$ , differing by a tower of height  $h > 1$  for which every cell is next to a boundary segment  $\ell$ . In one extension, this tower will be a valid tower move; in the other, it will not. Without loss of generality, suppose this tower stretches down from  $(\bar{x}, \bar{y})$  to  $(\bar{x}, \bar{y} - h + 1)$ , with  $\sigma(\bar{x}, \bar{y}) = 0$ ,  $\tau(\bar{x}, \bar{y}) = 1$ ,  $\sigma(\bar{x}, \bar{y} + h - 1) = c_1$  and  $\tau(\bar{x}, \bar{y} + h - 1) = c_1 + 1 \pmod{3}$ . See Figure 4.8, where this tower is of height 3 and  $c_1 = 2$ .

From  $\sigma$ , the only move that produces  $\tau$  requires picking the correct extension of  $\ell$ , the correct start location and color of the tower, and  $p < 1/2h$ :

$$P_C(\sigma, \tau) = \frac{1}{2} \cdot \frac{1}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{4h}.$$

Similarly, from  $\tau$ , the only move that produces  $\sigma$  requires picking the correct extension of  $\ell$ , the correct start location and color of the tower, and  $p < 1/2h$ :

$$P_C(\tau, \sigma) = \frac{1}{2} \cdot \frac{1}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{4h}.$$

We conclude that  $P_C(\sigma, \tau) = P_C(\tau, \sigma) = q/4h$ .

Let  $\sigma$  and  $\tau$  be 3-colorings of a free or mixed boundary region  $R$  where any reflex cells have fixed colors. Suppose  $\tau$  can be obtained from  $\sigma$  in one iteration of  $\mathcal{M}_C$  via a move that changes .

This completes our case analysis, proving that whenever  $\sigma$  and  $\tau$  differ by a move that changes cell(s) that touch most one segment of  $\partial R$ , then  $P_C(\sigma, \tau) = P_C(\tau, \sigma)$ .  $\square$

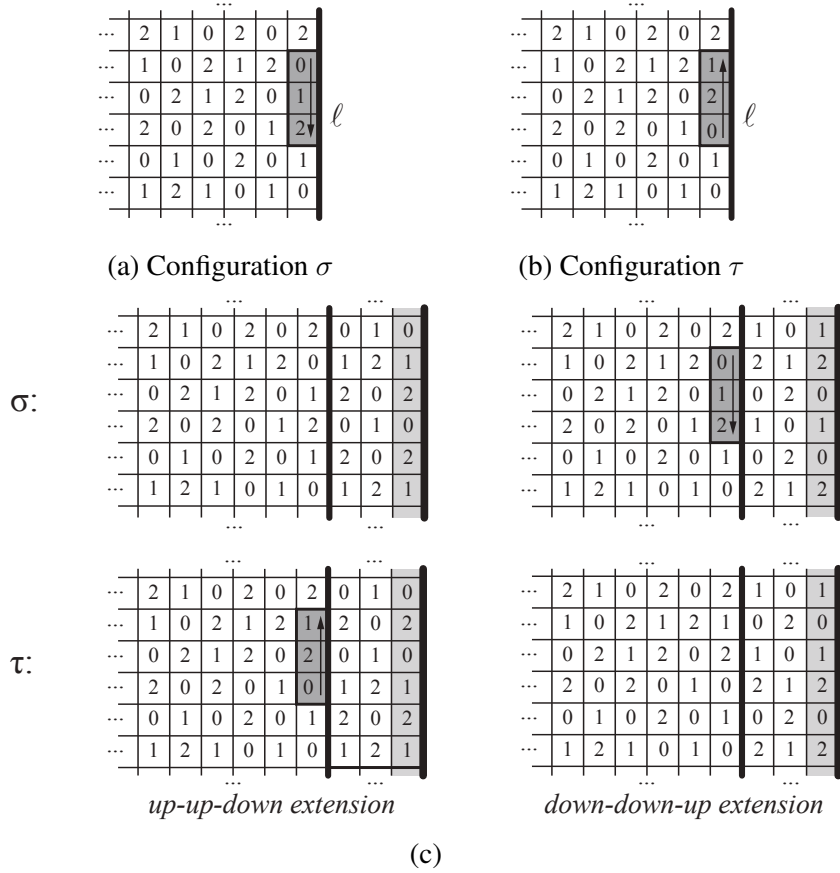


Figure 4.8: Colorings (a)  $\sigma$  and (b)  $\tau$  differ by a tower of height 3 (grey) adjacent to a boundary segment  $\ell$ . (c) The indicated tower moves for the random colorings of the extended region  $R'$  transition between  $\sigma$  and  $\tau$ .

This result will be critical to showing that  $\mathcal{M}_C$  converges to the uniform distribution, the distribution we'd like to be able to generate samples from.

#### 4.2.4 Where Two Boundary Segments Meet: Convex Corner

While the previous subsection describes how to extend region  $R$  beyond a single free boundary, in this section we begin to consider the relationships between extensions of different sides of  $R$ . In particular, we'll choose the extension of each side of  $R$  independently: the random extension chosen for one boundary segment of  $R$  will have no effect on the random extension chosen for a different boundary of  $R$ .

Care must be taken with regards to compatibility of random extensions of different




Figure 4.9: (a) A free boundary 3-coloring and the sides of one of its 16 extensions: the top and right sides are extended UUD, while the left and bottom sides are extended DDU. (b) The corners of the extension  $R'$  and the deterministic colors given to them based on the random extension chosen.

boundary segments of  $R$ . In the next section we discuss what to do when the random extensions of different sides of  $R$  overlap, as they would, for instance, when two boundary segments meet at a reflex corner. Here, we focus on what occurs when two boundary segments meet at a convex corner, with each having been extended independently and uniformly either UUD or DDU. At each convex corner we also include in  $R'$  the nine cells between the extensions of the two sides that meet at the corner (so that  $R'$  also has a convex corner). Given coloring  $\chi'$  on the sides of the extension, the coloring within this corner is completed deterministically. If the two sides meeting at a corner have opposite UUD and DDU configurations, four colors within the corner are uniquely determined because  $\chi'$  must be a valid 3-coloring. In other cases, the cell in this corner of  $R'$  that is closest to  $R$  should be given the unique color different from the color in  $R$ 's convex corner cell that still yields a valid 3-coloring. In both cases, the remaining colors in the corners will not affect any Markov chain moves intersecting  $R$ , so can be completed arbitrarily; we choose to complete them canonically by repeatedly giving each cell with at least two colored neighbors the lowest admissible color. See Figure 4.9 for one example of how the four convex corners of

the extension  $R'$  of a free boundary rectangle  $R$  are colored.

We now examine moves that occur in  $R'$  near a convex corner, and how the probabilities of these moves in  $R'$  translate into probabilities of moves when looking only within  $R$ . As above, the following lemma will be critical in showing that  $\mathcal{M}_C$  converges to the uniform distribution.

**Lemma 4.7.** *Let  $\sigma$  and  $\tau$  be 3-colorings of a free or mixed boundary region  $R$  where any reflex cells have fixed colors. Suppose  $\tau$  can be obtained from  $\sigma$  in one iteration of  $\mathcal{M}_C$  via a move that changes a convex corner cell. Then  $P_C(\sigma, \tau) = P_C(\tau, \sigma)$ .*

*Proof.* For such moves from  $\sigma$  to  $\tau$  that change the color of a cell at a convex corner where boundaries  $\ell_1$  and  $\ell_2$  meet, there are three cases to consider: Flips of the corner cell; tower moves that include one convex corner cell; and tower moves that include two convex corner cells, one at each end.

Case 1: Corner flip. Suppose  $\tau$  can be obtained from  $\sigma$  by flipping the color of a cell at a convex corner of  $R$ . Without loss of generality, suppose the boundary segments that meet at this corner are above and right of this corner cell  $(\bar{x}, \bar{y})$ ;  $\sigma(\bar{x}, \bar{y}) = 0$ ; and  $\tau(\bar{x}, \bar{y}) = 1$ . Adjacent cells  $(\bar{x} - 1, \bar{x})$  left of  $(\bar{x}, \bar{y})$  and  $(\bar{x}, \bar{y} - 1)$  below  $(\bar{x}, \bar{y})$  necessarily have color 2 in both tilings. There are four possible random extensions in the neighborhood of this corner cell  $(\bar{x}, \bar{y})$ , given by the particular extensions chosen (independently) for each of the two free boundaries that meet at  $(\bar{x}, \bar{y})$ . Across these four random extensions, there are 5 valid moves that transform  $\sigma$  into  $\tau$ , four towers whose intersection with  $R$  is of height 1 and one flip move; see Figure 4.10. Each occurs only if the correct of the four possible extensions, the correct start location, and the correct start color are chosen and, in the case of the tower moves, if  $p < 1/2$ . This means

$$P_C(\sigma, \tau) = \frac{1}{4} \cdot \frac{1}{3|R'|} \cdot \left(1 + 4 \cdot \frac{1}{2}\right) = \frac{3}{4}q.$$

It is simple to see, by permuting colors in the argument above, that there exist the same

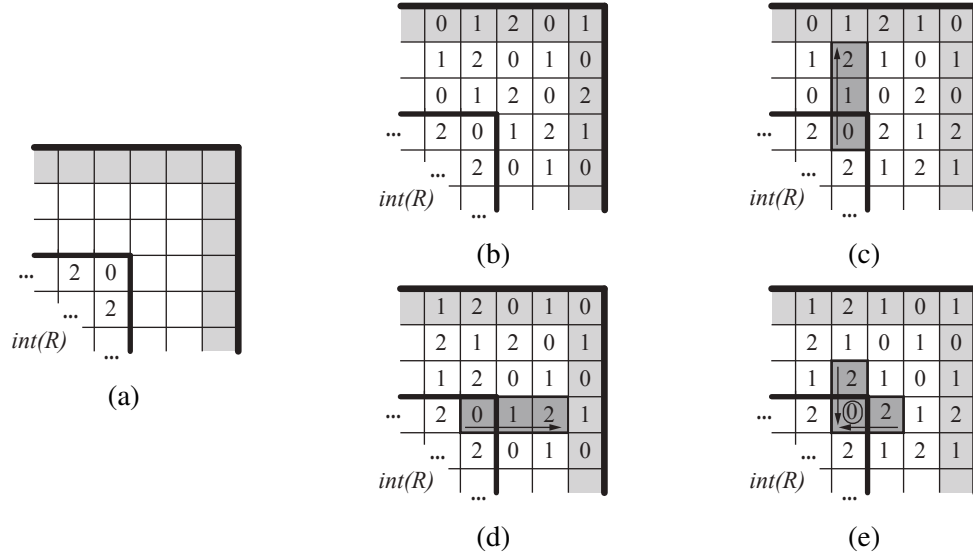


Figure 4.10: (a) A region  $R$  with a coloring  $\sigma$  that has a convex corner and the region  $R'$  we extend to near this convex corner. The four possible colorings of this extension in the neighborhood of this corner are shown in (b,c,d,e). Moves changing  $\sigma$  into  $\tau$ , which is the same as  $\sigma$  except for its corner color, are shown in dark gray; tower moves are indicated by arrows, and flips are circled.

number and types of moves transforming  $\tau$  into  $\sigma$ . We conclude that  $P_C(\sigma, \tau) = P_C(\tau, \sigma) = 3q/4$ .

Case 2: Tower including one convex corner cell. Suppose  $\sigma$  and  $\tau$  differ by a tower of height  $h \geq 2$  that contains a convex corner of  $R$ . Because reflex cells have fixed colors and cannot be in any tower move, this tower is adjacent to one of the two boundary segments that meet at the convex corner and abuts the other. Without loss of generality, suppose this tower stretches in the vertical direction from location  $(\bar{x}, \bar{y} - h + 1)$  up to corner location  $(\bar{x}, \bar{y})$ , the colors of the tower cells increase from bottom to top (mod 3),  $\sigma(\bar{x}, \bar{y}) = 1$ ,  $\tau(\bar{x}, \bar{y}) = 2$ ,  $\sigma(\bar{x}, \bar{y} + h - 1) = c_1$  and  $\tau(\bar{x}, \bar{y} + h - 1) = c_1 + 1 \pmod{3}$ . See Figures 4.11 and 4.12, where  $h = 3$  and  $c_1 = 0$ . Across the four possible random extensions at the corner, there are two moves for  $\sigma$  that yield  $\tau$ , and both are towers whose intersection with  $R$  is height 3. The same is true for moves for  $\tau$  yielding  $\sigma$ . Any such move requires picking the correct extension, the correct start location, the correct start color, and  $p < 1/2h$ . We

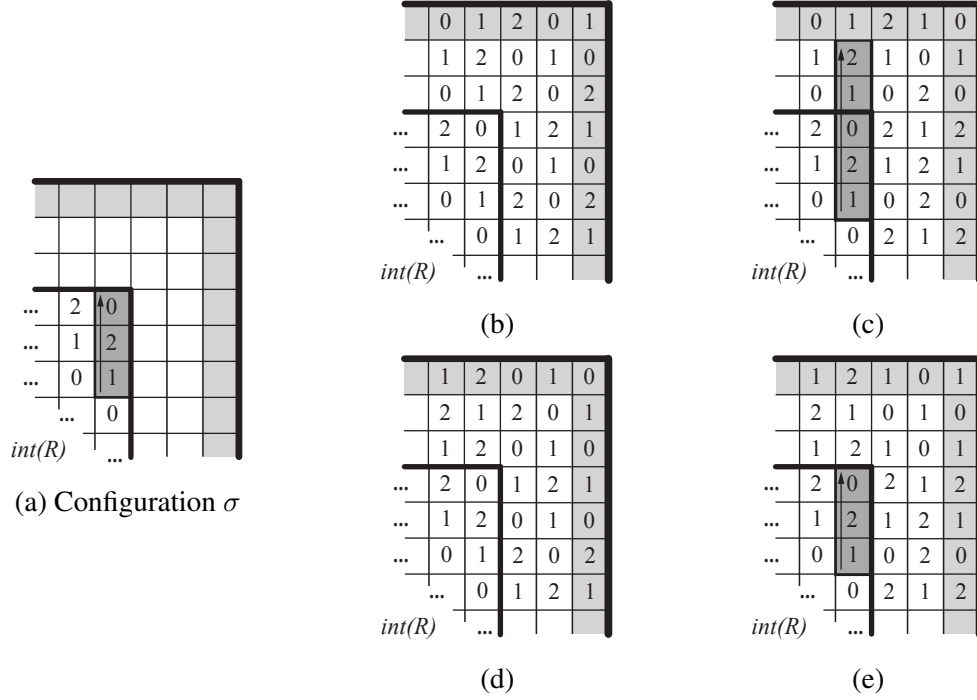


Figure 4.11: (a) A region  $R$  with a coloring  $\sigma$  that has a tower move including a convex corner of  $R$ , and the region  $R'$  we extend to near this convex corner. The four possible colorings of this extension in the neighborhood of this corner are shown in (b,c,d,e). Moves changing  $\sigma$  into  $\tau$  (Figure 4.12a) are shown in dark gray and indicated by arrows.

see that

$$P_C(\sigma, \tau) = P_C(\tau, \sigma) = 2 \cdot \frac{1}{4} \cdot \frac{1}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{4h}.$$

Case 3: Tower including two convex corner cells. We first note that it must be the case that the tower stretches the entire length of some boundary segment  $\ell_1$  of  $\partial R$ , meaning the start and end of the tower are the convex corner cells and every cell of the tower is adjacent to  $\ell_1$ ; this is because all reflex cells have fixed colors and cannot be in any valid tower move. Without loss of generality, we suppose this tower stretches vertically in  $\sigma$  from start cell  $(\bar{x}, \bar{y})$  up to end cell  $(\bar{x}, \bar{y} + h - 1)$ . Furthermore, suppose the colors increase (modulo 3) from bottom to top, with  $\sigma(\bar{x}, \bar{y}) = 0$ ,  $\tau(\bar{x}, \bar{y}) = 1$ ,  $\sigma(\bar{x}, \bar{y} + h - 1) = c_1$  and  $\tau(\bar{x}, \bar{y} + h - 1) = c_1 + 1(\text{mod } 3)$ . Just as in the previous case, for one extension of  $\ell_1$  tower moves from  $\sigma$  to  $\tau$  are possible, and in the other extension of  $\ell_1$  they are not. Across the two possible random extensions of each of the two other boundaries touching this tower at



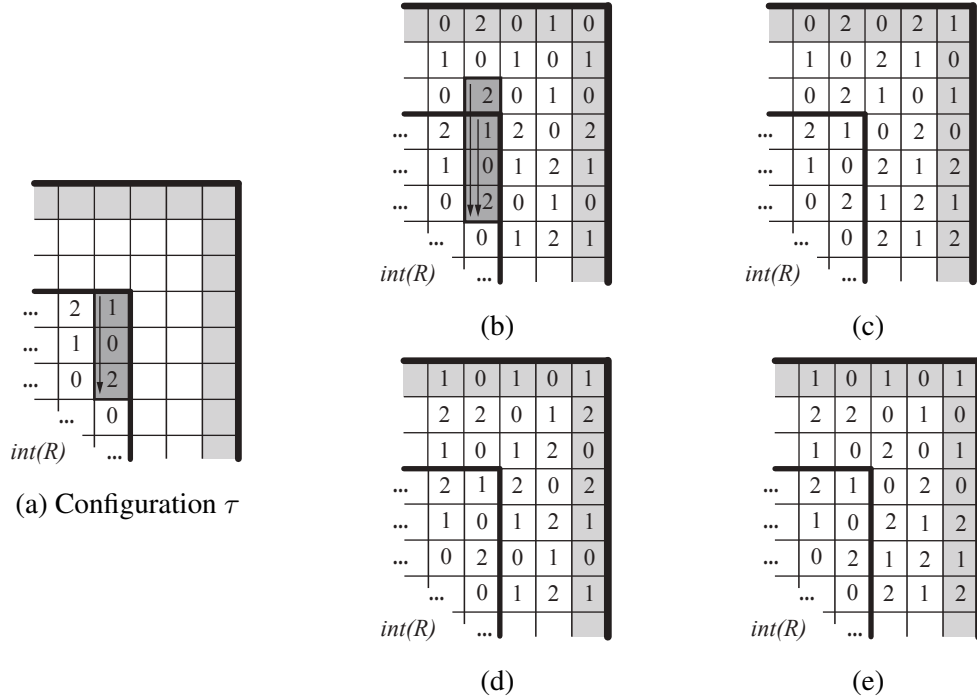


Figure 4.12: (a) A region  $R$  with a coloring  $\tau$  that has a tower move including a convex corner of  $R$ , and the region  $R'$  we extend to near this convex corner. The four possible colorings of this extension in the neighborhood of this corner are shown in (b,c,d,e). Moves changing  $\tau$  into  $\sigma$  (Figure 4.11a) are shown in dark gray and indicated by arrows.

its two ends, there are four moves for  $\sigma$  that yield  $\tau$ , and all are towers whose intersection with  $R$  is height  $h$ ; the moves locally near each end of the tower are just as in Figures 4.11 and 4.12. The same is true for moves for  $\tau$  yielding  $\sigma$ . Any such moves require picking the correct extension for each of the three boundaries touching this tower, the correct start location, the correct start color, and  $p < 1/2h$ . We see that

$$P_C(\sigma, \tau) = P_C(\tau, \sigma) = 4 \cdot \frac{1}{8} \cdot \frac{1}{3|R'|} \cdot \frac{1}{2h} = \frac{q}{4h}.$$

This concludes our proof that when  $\sigma$  and  $\tau$  differ by a move that changes the color of a convex corner cell,  $P_C(\sigma, \tau) = P_C(\tau, \sigma) = q/4h$ . □

This concludes our exploration of random extensions near convex corners, and we now move on to our final consideration: random extensions near reflex corners.

#### 4.2.5 Where Two Boundary Segments Meet: Reflex Corner

When two boundary segments of  $R$ , in particular boundary segments with free colors, meet at a reflex vertex, the extensions of both of these segments will overlap. It is unclear how one might resolve or otherwise consistently merge the extensions of both segments. We deal with this by considering combinatorial extensions, rather than planar extensions, and our requirement that every reflex cell has a fixed color plays a crucial role.

For a boundary segment ending at a reflex corner, we include the (fixed color) reflex cell as part of the boundary when making the extension. Extending this fixed reflex cell gives a fixed boundary on the side of the extension, ensuring the larger region  $R'$  is indeed a fixed boundary region. Extending reflex cells does cause overlaps with other cells of  $R$ , however.

In the case of reflex corners, our extended region  $R'$ , while locally planar, will not have a global embedding into the plane. When extensions of different boundary segments overlap, they are simply considered as different layers and they do not interact. We fix reflex cells, whose neighborhood could be made locally-nonplanar by such overlapping extensions, to ensure that the neighborhood of any cell without a fixed color is locally planar. See Figure 4.13, which shows how two free boundaries, which meet at a fixed color reflex cell, are independently extended in this way; these extensions overlap, meaning  $R'$  is not planar, but locally near every free color cell planarity is preserved.

This combinatorial approach conceptually enables random extensions near (fixed color) reflex corners, but would be challenging to actually implement due to its non-planarity. However, in proving reversibility for flips and tower moves near free boundaries (Lemmas 4.6 and 4.7), we've already calculated the probabilities of all possible moves when restricting our view to the original region  $R$ ; see the first two columns of Table 4.1. Using these probabilities, we can rewrite  $\mathcal{M}_C$  in a simpler way: instead of generating random extensions in each iteration, we make moves within  $R$  with the probabilities with which these moves would occur if we did generate random extensions. This is given as Algorithm 5.




Figure 4.13: (a) Two boundary segments of  $R$  meeting at a reflex corner, where the reflex cell (dark grey) has fixed color 0. (b) The extension of the vertical boundary segment, which includes extending the reflex cell. (c) The extension of the horizontal boundary segment, which includes extending the reflex cell. The extensions in (b) and (c) overlap, but are locally planar in the neighborhood of any free color cells.

---

**Algorithm 5** Markov Chain  $\mathcal{M}_C$  on grid region  $R$  with mixed boundary conditions

---

Starting at any initial 3-coloring  $\chi_0$  of  $R$ , repeat:

- 1: Choose, uniformly at random, a cell  $a$  of  $R$  that does not have a fixed color, a color  $c \in \{0, 1, 2\}$ , and a value  $p \in (0, 1)$ .
  - 2: **if**  $\chi_i(a) \neq c$  and no neighbors of  $a$  have color  $c$  in  $\chi_i$  **then**
  - 3:     **if**  $a$  is a convex corner of  $R$  **then**
  - 4:         Recolor  $a$  with color  $c$  with probability  $3/4$ .
  - 5:     **else** Recolor  $a$  with color  $c$  with probability 1.
  - 6: **else if**  $a$  starts a tower with start color  $c$ , height  $h \geq 2$ , and no fixed color cells **then**
  - 7:     **if** All cells of the tower are adjacent to a common boundary segment of  $R$  **then**
  - 8:         Make the tower move if  $p < 1/4h$ .
  - 9:     **else** Make the tower move if  $p < 1/2h$ .
- 

When looking at the changes made to the 3-coloring of  $R$  in each iteration, this algorithm varies from  $\mathcal{M}_C$  as stated in Algorithm 4 only in the probability of proposing a move: now we pick a cell of  $R$  with probability  $1/|R|$  instead of  $1/|R'|$ , where  $|R|$  is the number of cells in  $R$  without a fixed color. This has no effect on the asymptotic behavior of the chain, and in fact speeds it up. Additionally, Algorithm 5 doesn't need to generate random extensions in each iteration, providing even more computational savings. It is important to emphasize that while we've simplified  $\mathcal{M}_C$  into a Markov chain that makes updates to colors only within  $R$ , without considering random extensions at all, it was our approach of random extensions that allowed us to derive and verify the probabilities with which the moves of Algorithm 5 should be made. Furthermore, it was our insights into randomized

Type of move	Prob. in $\mathcal{M}_C$	Prob. in $\widetilde{\mathcal{M}}$ [61]
Interior flip	$q$	$s$
Boundary flip	$q$	$s$
Corner flip	$\frac{3q}{4}$	$\frac{3s}{4}$
Interior tower, height $h > 1$	$\frac{q}{2h}$	$\frac{s}{2h}$
Tower abutting one boundary, height $h > 1$	$\frac{q}{2h}$	$\frac{s}{2h}$
Tower abutting two boundaries, height $h > 1$	$\frac{q}{2h}$	$\frac{s}{2h}$
Tower adjacent to one boundary, height $h > 1$	$\frac{q}{4h}$	$\frac{s}{4h}$
Tower including one corner cell, height $h > 1$	$\frac{q}{4h}$	$\frac{s}{4h}$
Tower including two corner cells, height $h > 1$	$\frac{q}{4h}$	$\frac{s}{4h}$

Table 4.1: Types of moves for regions with some free boundaries, and the probability with which they occur in  $\mathcal{M}_C$  and, specifically for rectangular regions, in  $\widetilde{\mathcal{M}}$  from [61]. Here  $q = \frac{1}{3|R|}$  and  $s = \frac{1}{12mn}$ .

extensions that allowed us to realize sampling from mixed boundary regions was possible, determine what the conditions on non-convex regions must be to allow such sampling to be efficient, and decide the probabilities of different types of moves. Despite this, it is the formulation of  $\mathcal{M}_C$  given in Algorithm 5 that we will use throughout the rest of this section.

When specifically considering free boundary rectangles, Algorithm 5 is identical to Algorithm  $\widetilde{\mathcal{M}}$  of [61], up to the probability of proposing a move (that is, of picking a cell and a color); see the third column of 4.1. In  $\widetilde{\mathcal{M}}$ , in an  $m \times n$  rectangle a move is proposed with probability  $1/(12mn)$ ; in Algorithm 5, that probability would be  $1/(3mn)$ .<sup>1</sup>

#### 4.2.6 Properties of Markov Chain $\mathcal{M}_C$ .

Now that we have defined Markov chain  $\mathcal{M}_C$ , we can discuss its properties.

---

<sup>1</sup>Algorithm  $\widetilde{\mathcal{M}}$  of [61] has a stationary probability of  $3/4$  in each iteration for technical reasons, to ensure that all eigenvalues of the transition matrix are at least  $1/2$ , which is helpful when applying the comparison technique [43] to obtain mixing results for Glauber dynamics.

**Lemma 4.8.** *Markov chain  $\mathcal{M}_C$  is aperiodic.*

*Proof.* Each state  $\sigma$  has stationary probability  $P_C(\sigma, \sigma) \geq 1/3$ ; whenever the color  $c$  picked is the same as the color of the cell selected, no move occurs.  $\square$

Unlike most Markov chains considered in this thesis,  $\mathcal{M}_C$  may not be irreducible, and for many regions  $R$  with mixed boundary conditions it is not. However, whenever  $\mathcal{M}_C$  is not irreducible, we will still be able to study its behavior on each piece of its state space. We now characterize on which state spaces  $\mathcal{M}_C$  is irreducible by appealing to height functions of tilings (Section 4.1.5).

**Lemma 4.9.** *Let  $\Omega$  be the set of all 3-colorings of some grid region  $R$  with mixed boundary conditions where at least one cell has a fixed color and all reflex cells have a fixed color. Pick any fixed boundary cell  $a$ , and for each configuration  $\sigma \in \Omega$ , define  $h_\sigma(\cdot)$  to be the unique height function for  $\sigma$  that has  $h_\sigma(a) = i$ . Let  $\bar{\Omega}$  be a maximal subset of  $\Omega$  such that for all fixed boundary cells  $b$ ,  $h_\sigma(b)$  is constant for all  $\sigma \in \bar{\Omega}$ . Then,  $\mathcal{M}$  is ergodic on  $\bar{\Omega}$ .*

*Proof.* It is straightforward to show, using a standard height function argument, that if  $\sigma$  and  $\tau$  are both in some  $\bar{\Omega} \subseteq \Omega$ , meaning for all fixed boundary cells  $b$  that  $h_\sigma(b) = h_\tau(b)$ , then there exists a sequence of flip moves that transforms  $\sigma$  into  $\tau$  and vice versa. Because this is the same approach that has already been used to show irreducibility of Markov chains on fixed and free boundary 3-colorings in both [79] and [61], we do not reproduce the details here. It is precisely the fact that every 3-coloring in  $\bar{\Omega}$  has the same height values at all fixed boundary locations that enables this proof to go through; flip moves and tower moves are insufficient to change the height of fixed color cells.

Because flip moves are a subset of the moves of  $\mathcal{M}_C$ , this is enough to show  $\mathcal{M}_C$  is irreducible on  $\bar{\Omega}$ . Because  $\mathcal{M}_C$  is aperiodic (Lemma 4.8), it is also ergodic on  $\bar{\Omega}$ , completing the proof.  $\square$

This lemma addresses all of our regions of interest except those without any fixed color cells. Because we assume that any reflex cells have fixed colors, the only remaining case is

free boundary regions without any reflex cells, which are precisely free boundary rectangles, as considered in [61]. Their work gives us the following result.

**Lemma 4.10** ([61]). *Let  $\Omega$  be the set of all 3-colorings of some grid rectangle  $R$  with free boundary conditions. Then  $\mathcal{M}_C$  is ergodic on  $\Omega$ .*

*Proof.* Our chain  $\mathcal{M}_C$  and tower Markov chain  $\widetilde{\mathcal{M}}$  of [61] are identical up to the probability of proposing a move. Because  $\widetilde{\mathcal{M}}$  is ergodic, so is  $\mathcal{M}_C$ .  $\square$

For mixed boundary conditions, in the special case where all  $\sigma \in \Omega$  map to height functions  $h_\sigma$  that are constant on all fixed boundary cells (where  $\overline{\Omega} = \Omega$ ), recall that these mixed boundary conditions are *height consistent*. We note that height consistency can depend on the values of the fixed boundary cells; the example in Figure 4.4b is not height consistent, but if one of its fixed cells had a different value it would be.

**Corollary 4.11.** *Markov chain  $\mathcal{M}_C$  on a grid region  $R$  with height-consistent mixed boundary conditions and all reflex cells fixed is ergodic.*

*Proof.* This follows immediately from Lemma 4.9 when  $R$  has at least one fixed cell, because in this case  $\Omega = \overline{\Omega}$ . When no boundary cells have fixed colors, the region is vacuously height-consistent and must be a rectangle, and ergodicity follows from Lemma 4.10.  $\square$

**Lemma 4.12.** *If  $\mathcal{M}_C$  on a mixed boundary grid region  $R$  with all reflex cells fixed is ergodic on some state space  $\overline{\Omega} \subseteq \Omega$ , then it converges to the uniform distribution over  $\overline{\Omega}$ .*

*Proof.* Because  $\mathcal{M}_C$  is ergodic and finite, it converges to a unique stationary distribution. By Lemmas 4.6 and 4.7, for any  $\sigma, \tau$  in  $\Omega$  that differ by a single move – a flip or a tower, interior or touching boundary segment(s) – we have that  $P_C(\sigma, \tau) = P_C(\tau, \sigma)$ . No other types of moves exist, other than those considered in these two lemmas, because reflex cells have fixed colors and so cannot be part of any move. We can then verify that the uniform distribution over  $\overline{\Omega}$  is  $\mathcal{M}_C$ 's stationary distribution using detailed balance:

$$\pi(\sigma)P(\sigma, \tau) = \frac{1}{|\overline{\Omega}|}P(\sigma, \tau) = \frac{1}{|\overline{\Omega}|}P(\tau, \sigma) = \pi(\tau)P(\tau, \sigma).$$

□

**Corollary 4.13.** *Markov chain  $\mathcal{M}_C$  on a mixed boundary grid region  $R$  with height consistent boundary conditions and all reflex cells fixed converges to the uniform distribution over  $\Omega$ .*

*Proof.*  $\mathcal{M}_C$  is ergodic on  $\Omega$  by Corollary 4.11, so this follows immediately from Lemma 4.12.

□

#### 4.2.7 Mixing Time of Markov Chain $\mathcal{M}_C$ .

Our final consideration for  $\mathcal{M}_C$  is its convergence time. Because of our randomized extensions approach, a path coupling argument showing fast mixing comes almost for free: for a tower Markov chain on fixed boundary 3-colorings, previous work has already shown that if two configurations differ by a single cell then in expectation they get no farther apart after one iteration [79].  $R'$  is a fixed boundary region, and after generating the random extension we are almost exactly making a move of the fixed boundary tower-based chain on  $R'$ . Based on this, one might expect that two configurations differing at a single location in expectation get no farther apart in each random extension, and thus get no farther apart overall in one iteration of  $\mathcal{M}_C$ . Two small considerations must be taken into effect: tower moves in  $R'$  are made with probability  $1/2h$ , where  $h$  is the size of the tower's intersection with  $R$ , not its overall height; and when two configurations differ at a boundary cell, this difference permeates into the random extension as well, meaning the colorings of  $R'$  no longer differ only at a single cell.

However, despite these concerns, we are able to adapt the path coupling argument of [79] to prove that  $\mathcal{M}_C$  is rapidly mixing whenever it is ergodic. A full proof of fast mixing for  $\mathcal{M}_C$  as stated in Algorithm 4, via path coupling across all possible random extensions, can be found in [25]. Here, for the sake of simplicity and clarity, we prove fast mixing for the essentially equivalent statement of  $\mathcal{M}_C$  as Algorithm 5, which only

looks within  $R$ . This approach was also used by [61] to show rapid mixing of their tower-based Markov chain on free boundary 3-colorings, and instead of proving fast mixing from scratch we use some of their intermediate results.

We use a path coupling argument (Theorem 2.4) to show that whenever  $\mathcal{M}_C$  is ergodic, it is rapidly mixing. Consider a joint process  $(\mathcal{A}, \mathcal{B})$  on  $\Omega \times \Omega$ , where each of  $\mathcal{A}$  and  $\mathcal{B}$  is a copy of Markov chain  $\mathcal{M}_C$ . Let  $A_t$  and  $B_t$ , respectively, be their marginal distributions at iteration  $t$ . We couple by making the same choice of cell  $a \in R$ , color  $c \in \{0, 1, 2\}$ , and probability  $p \in (0, 1)$  for both  $\mathcal{A}$  and  $\mathcal{B}$  in Step 1 in each iteration of Algorithm 5.

We define the distance  $\Phi$  between two 3-colorings of  $R$  to be the minimum number of flips needed to transform one 3-coloring into the other. Our goal is to show, over time, that  $A_t$  and  $B_t$  get closer together in expectation according to metric  $\Phi$ . Because we use the path coupling approach, we only have to consider the change in  $\Phi$  after one iteration for  $A_t$  and  $B_t$  differing at exactly one cell.

In the next two lemmas, we consider this case where  $A_t$  and  $B_t$  differ by a single flip, and we analyze the change in distance  $\Delta\Phi_t := \Phi(A_{t+1}, B_{t+1}) - \Phi(A_t, B_t)$ . Because of the way  $\mathcal{M}_C$  can be simply stated in Algorithm 5, we are able to leverage previous work in [61] to prove our results without significant additional effort.

**Lemma 4.14.** *If marginal 3-colorings  $A_t$  and  $B_t$  differ by a single flip, then  $\mathbb{E}[\Delta\Phi_t] \leq 0$ .*

*Proof.* Marginal 3-colorings  $A_t$  and  $B_t$  must differ by an interior flip, a boundary flip, or a convex corner flip; because of the assumptions that we make about region  $R$  (Section 4.1.3), these are the only possible cases.

Coalescence of a coupling of a tower Markov chain where the marginal 3-colorings differ by interior flips was analyzed in [79]; in  $\mathcal{M}_C$ , all moves coalescing two chains differing by a single flip occur with the same probabilities as in [79] (up to the probability of proposing a move), but moves moving the two marginal chains farther apart may happen with less probability, for instance if a tower near the interior flip is adjacent to a boundary or includes fixed cells. We conclude that for interior flips,  $\mathbb{E}[\Delta\Phi_t] \leq 0$ .



Coalescence of a coupling where the marginal 3-colorings differ by boundary and corner flips was analyzed in [61] for a tower Markov chain on free boundary rectangles, which uses the same probabilities that we've derived (up to the probability of proposing a move). Even though we consider a much larger class of regions, locally any flip next to a free boundary looks the same regardless of whether we're considering a free boundary rectangle or some other mixed boundary region, up to the presence of fixed color cells. However, the presence of fixed color cells does not decrease the probability that two chains differing by a boundary or corner flip coalesce, it only potentially decreases the probability that the two chains move farther apart. We thus conclude, based on the work of [61], that in this case as well  $\mathbb{E}[\Delta\Phi_t] \leq 0$ .

As previously mentioned, a proof of this result by appealing directly to randomized extensions can be found in [25]; we reuse previous work here for the sake of simplicity, clarity, and length.  $\square$

**Lemma 4.15.** *For marginal 3-colorings  $A_t$  and  $B_t$ , it holds that*

$$\mathbb{P}(\Delta\Phi_t \neq 0) \geq \frac{1}{6|R|^2}.$$

*Proof.* We must prove this lemma for all possible  $A_t$  and  $B_t$ , not just those differing by a flip move, because, unlike expected changes in distance, variances cannot be added along paths. Since a proof of this fact would be nearly identical to that of the same result in [61], we do not include the detailed case analysis. Following the exact argument in Section 6.2 of [61] and using the fact that the length of any tower is at most  $|R|$  gives the bound in the lemma.  $\square$

Using the above lemmas, we now show that whenever  $\mathcal{M}_C$  is ergodic, it mixes in polynomial time.

**Theorem 4.16.** *If Markov chain  $\mathcal{M}_C$  on a mixed boundary region  $R$  where all reflex cells have fixed colors is ergodic on a state space  $\overline{\Omega} \subseteq \Omega$ , it mixes over this state space in*

polynomial time, at most  $O(|R|^6)$  where  $|R|$  is the number of cells in  $R$  without fixed colors.

*Proof.* We apply the path coupling theorem (Theorem 2.4) using the same distance function  $\Phi$  as above: for two configurations  $\sigma$  and  $\tau$ ,  $\Phi(\sigma, \tau)$  is the minimum number of flips necessary to transform  $\sigma$  into  $\tau$ .

To verify the hypotheses of Theorem 2.4, we first note that metric  $\Phi$  takes on integer values in  $[0, |R|^2]$  (Lemmas 4.4 and 4.5). Metric  $\Phi$  also satisfies the stated path condition: for  $U$  the set of all pairs of colorings differing by a single flip, the distance between any two configurations is exactly the shortest path distance between them along edges in  $U$ . If  $A_t$  and  $B_t$  differ on a single flip, then by Lemma 4.14 we have  $\mathbb{E}[\Delta\Phi_t | A_t, B_t] \leq 0$ . Furthermore, by Lemma 4.15, we see that  $\mathbb{P}(\Delta\Phi_t \neq 0) \geq 1/6|R|^2$ . By Theorem 2.4, where  $S = |R|^2$  and  $\alpha = 1/6|R|^2$ , we conclude that the mixing time of  $\mathcal{M}_C$  satisfies

$$t_{mix}(\varepsilon) \leq \left\lceil \frac{eS^2}{\alpha} \right\rceil \lceil \log(\varepsilon^{-1}) \rceil \leq O(|R|^6 \log(\varepsilon^{-1})).$$

□

**Theorem 4.17.** *Markov chain  $\mathcal{M}_C$ , on a region  $R$  that has height-consistent mixed boundary conditions and all reflex cells fixed, mixes in polynomial time at most  $O(|R|^6)$ .*

*Proof.* By Corollary 4.11,  $\mathcal{M}_C$  is ergodic on the state space  $\Omega$  of all valid 3-colorings of  $R$  subject to its height consistent mixed boundary constraints. Applying Theorem 4.16 then completes the proof. □

**Corollary 4.18.** *If Glauber dynamics on a mixed boundary region  $R$  where all reflex cells have fixed colors is ergodic on a state space  $\bar{\Omega} \subseteq \Omega$ , it mixes over this state space in polynomial time.*

*Proof.* This follows directly from the comparison argument in [97] and is very similar to the comparison argument given for free-boundary 3-colorings in [61]. □

It is also possible to use  $\mathcal{M}_C$  to sample from  $L$ -shaped regions with free boundary conditions. This is the first such sampling result for any non-rectangular region with free boundary conditions.

**Corollary 4.19.** *Let consider grid region  $R$  with exactly one reflex cell and free boundary conditions. It is possible, in polynomial time, to generate a uniformly random sample for  $R$ .*

*Proof.* Uniformly at random fix the color of the cell at the reflex corner of  $R$ ; in a uniformly random 3-coloring of  $R$ , this cell is equally likely to be any of the three colors. Now,  $R$  is a region with height-consistent mixed boundary conditions where all reflex cells are fixed, and by Theorem 4.17 we can use  $\mathcal{M}_C$  to generate a random sample from  $R$  in polynomial time. □

The above results only allow efficient sampling when any mixed boundary conditions are height consistent, which is a restrictive and slightly unnatural set. In the next subsection we appeal to the dichotomy between approximate sampling and approximate counting, which will enable us to push past this barrier.

#### 4.2.8 Self-reducibility and Approximate Counting with $\mathcal{M}_C$ .

We have thus far focused only on sampling. Seminal work of Jerrum, Valiant and Vazirani [71] shows how efficient algorithms for sampling can be used to construct efficient algorithms to approximately count, provided the underlying problem is *self-reducible*. Algorithms that require regions to have a specific shape, such as a rectangle in  $\mathbb{Z}^2$ , typically fail to be self-reducible. The reduction requires incrementally fixing parts of a configuration and sampling from the remainder subject to these constraints. For 3-colorings, this means incrementally fixing cells to have their most likely color and sampling from nearly rectangular regions where the colors of some of the boundary cells have been fixed. Figure 4.14 gives a typical step of the reduction from sampling to approximate counting. Our




Figure 4.14: (a) An intermediate step of the reduction between sampling and approximate counting, where the colors of some cells of a free boundary rectangle have been fixed. (b) The resulting mixed boundary region we sample from in one step in the approximate counting process.

mixed boundary Markov chains allow us to sample from exactly those regions that appear in the reduction from sampling to counting. Any reflex corners on the boundary will be fixed, and all fixed boundary cells will be contiguous. This means Theorem 4.16 applies and  $\mathcal{M}_C$  can be used to sample from such regions, and we can use such samples iteratively to approximately count.

This approach based on random extensions yields the first efficient algorithm for approximately counting 3-colorings on rectangles with free boundaries. This is possible precisely because sampling from the set of height-consistent mixed boundary regions is exactly what is needed for self-reducibility. In fact, we can approximately count the number of 3-colorings of any grid region with height consistent mixed boundary constraints and all reflex cells fixed.

#### 4.2.9 Using Approximate Counts to Sample from More Regions.

Because sampling from height-consistent mixed boundary regions enables approximate counting, it is possible to use this approximate counting algorithm as a subroutine to get an efficient sampling process for a much larger class of mixed boundary and free boundary regions. For non-height consistent mixed-boundary regions, the only barrier to efficient sampling is that the state space isn't connected;  $\mathcal{M}_C$  still converges to the uniform distribu-

tion and mixes rapidly on each component of  $\Omega$ . Counting enables us to know the relative sizes of the components of  $\Omega$ , and thus known the probability at uniformity of being in a particular component. To generate a uniform sample over  $\Omega$ , we can pick a component of  $\Omega$  with probability proportional to its size, and then generate a uniform sample within that component.

This gives an efficient sampling algorithm whenever the number of components in  $\Omega$  is polynomial. For example, this includes any regions with a constant number of fixed boundary components, or a logarithmic number of fixed boundary components that can be joined by constant-length paths. For free boundary regions with multiple reflex corners, as long as there are at most a polynomial number of ways to set the relative heights at the reflex corners, the same approach enables efficient sampling.

It is clear that, while being able to sample from height-consistent mixed boundary conditions is only a modest improvement over previous work, it includes exactly the cases necessary to enable much broader sampling algorithms via approximate counting.

### 4.3 Lozenge Tilings

As was the case for 3-colorings, sampling lozenge tilings on the triangular lattice has been known to be possible for fixed boundary regions [79], but techniques do not extend to sampling for free boundaries. A *lozenge tiling* is a covering of a region of the triangular lattice with rhombus shaped lozenges, each covering exactly two adjacent triangles, so that every triangle is covered by a unique lozenge (it is dual to perfect matchings on hexagonal lattice regions). This 2-dimensional tiling problem has received a lot of attention because of its remarkable properties: each lozenge tiling seems to jump out of the page to form a 3-dimensional structure comprised of supported boxes, and indeed this “height function” has allowed many deep mathematical discoveries, most notably the “Arctic circle theorem” [31].

Luby et al. [79] showed a tower chain mixes in polynomial time for lozenge tilings in




Figure 4.15: (a) A free boundary lozenge tiling of a triangular region  $E$ , and (b - c) the two possible random extensions for its top-right side.

any simply connected region with fixed boundary conditions, where lozenges are required to remain entirely within the region. The comparison argument of Randall and Tetali [97] again shows Glauber dynamics also converge quickly. There has also been some interest in the free boundary case, where lozenges may overlap the boundary arbitrarily (see Figure 4.15a). This was studied by Martin and Randall [82] using a correspondence between tilings and non-intersecting lattice paths, dynamic programming, and an approach based on determinants. Their results include an algorithm for approximately counting the number of lozenge tilings of hexagonal regions with free boundaries, but their method does not seem to generalize to other regions so we cannot, for instance, use self-reducibility to construct an algorithm for efficient sampling.

Just as for 3-colorings, we can define *randomized extensions* for lozenge tilings; the random extensions for one side of a free boundary lozenge tilings are shown in Figures 4.15b and 4.15c, and one completed random extension is shown in Figure 4.16. Using these randomized extensions we can efficiently sample from a larger class of free and mixed boundary lozenge tilings than was previously possible. This new class includes precisely those mixed boundary conditions we need for self-reducibility, which we can use to approximately count the free boundary lozenge tilings of triangular regions. In turn, this ability to approximately count certain lozenge tilings can be used to sample from even more regions, as above. More details about lozenge tilings were omitted due to length but can be




Figure 4.16: A random extension, including all sides and completed corners, of a free boundary lozenge tiling of an equilateral triangle  $E$ .

found in [25]. Such results about lozenge tilings provide evidence that this technique of randomized extensions is not limited to the specific case of grid 3-colorings but is in fact more generally applicable.

## CHAPTER 5

### MARKOV CHAIN ALGORITHMS FOR PROGRAMMABLE MATTER

In this section we consider algorithmic foundations for *programmable matter*, a material or substance that utilizes user input or stimuli from its environment to change its physical properties in a programmable fashion. We are motivated by a desire to better understand the capabilities and limitations of such systems, where simple components collectively self-organize to solve system-wide problems in a decentralized way. Towards this goal, rather than focusing on specific instances of programmable matter, we instead consider an abstraction, *self-organizing particle systems*, that captures many features common across different active programmable matter systems. This abstraction allows rigorous exploration of general algorithmic capabilities and limitations free from the particular nuances and constraints of certain instantiations.

In this chapter, we present a new method to using Markov chains to develop robust distributed algorithms for self-organizing particle systems. By harnessing the interplay between local moves and emergent behavior of Markov chains, we can achieve a variety of desired objectives, including *compression*, described in Sections 5.3-5.5, and *shortcut bridging*, described in Sections 5.6-5.8.

#### 5.1 The Stochastic Approach to Self-Organizing Particle Systems

In a *self-organizing particle system*, a collection of simple computational elements (to be referred to as *particles*) with limited computational power each perform fully distributed, local, asynchronous algorithms to solve system-wide problems such as movement, configuration, and coordination. We focus on the *geometric amoebot model* [39], where particles occupy the vertices and move along the edges of the infinite triangular grid graph  $\Gamma$ . We are interested in what these simple particles can collectively accomplish without any cen-



tralized control.

In [21], we initiated the *stochastic approach* to developing programmable matter algorithms. By harnessing our knowledge of the relationship between local moves of Markov chains and their emergent behavior, we can accomplish remarkable objectives with distributed algorithms. The motivation underlying the use of Markov chains to accomplish objectives in programmable matter comes from statistical physics, where ensembles of particles represent physical systems and demonstrate that local micro-behavior can induce global macro-scale changes to the system [7, 12, 100]. Like a spring relaxing, physical systems favor configurations that minimize energy. Each configuration  $\sigma$  has energy determined by a *Hamiltonian*  $H(\sigma)$  and weight  $w(\sigma) = e^{-B \cdot H(\sigma)} / Z$ , where  $B = 1/T$  is inverse temperature and  $Z = \sum_{\tau} e^{-B \cdot H(\tau)}$  is the normalizing constant known as the *partition function*. Markov chains have been well-studied as a tool for sampling configurations of these systems proportional to their energy function  $w(\sigma)$ , where the configurations with the lowest values of  $H(\sigma)$  are most likely to be sampled.

In the stochastic approach to programmable matter, we introduce a Hamiltonian  $H(\sigma)$  over particle configurations  $\sigma$  that assigns the lowest values to desirable configurations; we then design a Markov chain algorithm to favor these configurations with small Hamiltonians. For example, in Sections 5.3 to 5.5, we consider the *compression problem*, where desirable configurations have short perimeter. In this case, we use as a Hamiltonian  $H(\sigma) = p(\sigma)$ , the perimeter of  $\sigma$ . Given a Hamiltonian capturing our desired objective, and setting  $\lambda = e^B$ , we get  $w(\sigma) = \lambda^{-H(\sigma)}$ , where  $w(\sigma)$  is the relative likelihood of configuration  $\sigma$  in the stationary distribution of our Markov chain. As  $\lambda$  gets larger (by increasing  $B$ , effectively lowering temperature), we increasingly favor configurations where  $H(\sigma)$  is small.

By harnessing our knowledge of how the probabilities of local moves induce global emergent behavior, we can design a Markov chain  $\mathcal{M}$  that uses only local moves and eventually reaches a distribution that favors configurations proportional to their weight  $w(\sigma)$ ,

so particle configurations with smaller  $H(\sigma)$ , which have our desired property, are more likely. Rather than terminating the process at some point and using the configuration at that time step as a random sample, we instead run the Markov chain indefinitely, moving among different configurations but remaining at the stationary distribution of  $\mathcal{M}$ , which for the applications we consider we prove meets our desired objectives with high probability.

In each iteration  $\mathcal{M}$  changes the position of one particle by at most one unit, and only local information is used to decide with what probability to make each such move. Because of these careful considerations in the design of our Markov chain  $\mathcal{M}$ , it can be implemented in a distributed fashion by a particle system without centralized control, yielding a fully distributed, local, asynchronous algorithm, as desired. Because the distributed algorithms we develop for particle systems come from carefully-designed Markov chains, they have several nice features:

- we can use the many Markov chain analysis tools available to us to provide guarantees about their behavior;
- they are the first algorithms for self-organizing particle systems to be robust to a variety of types of particle failures; and
- the time-independence of Markov chains guarantees near-obliviousness, with particles only having to remember one bit of persistent memory over time.

The latter two features are particularly important for developing distributed algorithms that are implementable in real world settings, where failures are common and maintaining persistent memory is difficult [56].

We begin by formalizing the model we consider, and then present our Markov chain algorithms and their associated distributed implementations for two biologically-inspired problems: *compression* and *shortcut bridging*.




Figure 5.1: (a) A section of the triangular lattice  $\Gamma$ ; (b) expanded particles (each denoted by its two occupied adjacent locations in  $\Gamma$  and a thick line in between) and contracted particles (occupying one location); (c) two non-neighboring contracted particles with different offsets for their labels of the lattice edges leading to neighboring locations.

## 5.2 Background on Self-Organizing Particle Systems

We begin by describing our model, the geometric amoebot model for programmable matter and define some properties of particle systems.

### 5.2.1 The Geometric Amoebot Model

The Geometric Amoebot Model, inspired by the behavior of amoeba, was first proposed in 2014 [39] to model interacting computational particles, and has since served as the underlying model for a number of papers on algorithmic foundations of self-organizing particle systems and programmable matter. This includes work on problems such as shape formation [38], object coating [36, 41], and leader election [35].

In the *amoebot model*, programmable matter consists of particles whose properties we now detail. An infinite undirected graph  $G = (V, E)$  represents the set of relative locations that the particles can occupy (given by  $V$ ) and the set of all possible atomic transitions between locations in  $V$  (given by  $E$ ). We further assume the *geometric* variant of the amoebot model, which imposes an underlying geometric structure  $G = \Gamma$ , where  $\Gamma$  is the *triangular lattice* shown in Figure 5.1a (also called the infinite regular triangular grid graph, and denoted by  $G_{egt}$  in earlier work).

Each particle occupies either a single location (i.e., it is *contracted*) or a pair of two

adjacent locations on the graph (i.e., it is *expanded*); Figure 5.1b illustrates expanded and contracted particles on  $\Gamma$ . Each location can be occupied by at most one particle (that is, the particles follow an *exclusion process*). Particles achieve movement via a series of *expansions* and *contractions*: a contracted particle may expand into an adjacent unoccupied location to become expanded, and completes its movement by contracting to once again occupy only one location.

Two particles occupying adjacent nodes are said to be *neighbors*. Particles are *anonymous*, but can uniquely identify each one of their possible neighboring locations and check which of those neighboring locations are occupied by particles. We assume particles have a common *chirality*, meaning they share the same notion of *clockwise direction*, allowing them to label their neighboring locations in clockwise order.<sup>1</sup> However, particles do not share a global orientation and thus may have different offsets for their labels of neighboring locations (Figure 5.1c).

Every particle has a constant-size, shared, local memory which both it and its neighbors can read from and write to for communication. Because of the limitation on memory size, particles can know neither the total size of the system nor an estimate of it. Particles execute a sequence of *atomic actions*, in each of which they do some local computation (in our case, this may involve checking which of its adjacent locations are occupied with particles) and an expansion or contraction. We assume a fully asynchronous system, where particles perform atomic actions concurrently and at different, possibly variable speeds; conflicts, which in our context arise when two particles attempt to expand into the same location, are resolved in an arbitrary manner. In order to analyze such systems, we use the standard asynchronous model from distributed computing (see, e.g., [80]), allowing us to evaluate the progress of the system through a sequential series of individual *particle activations*, where every time a particle is activated, it performs an atomic action. A classical

---

<sup>1</sup>Formally, particles assign labels to the lattice edges leaving the location(s) they occupy, meaning an expanded particle assigns two labels to the each of the two locations adjacent to both its head and its tail. We do not require this level of detail in the results we present here.

result under this model states that for any concurrent asynchronous execution of atomic actions, there is a sequential ordering of actions producing the same end result, provided conflicts that arise in the concurrent execution are resolved. Thus, while in reality many particles may be active concurrently, it suffices when analyzing our algorithms to consider a sequence of activations where only one particle is active at a time.

### 5.2.2 Terminology for Particle Systems

We introduce notation and terminology that will be used throughout this paper. We call the collection of locations in  $\Gamma$  that are occupied by particles an *arrangement*; note two arrangements are the same even if different particles occupy the locations within the arrangement. We can define an equivalence relation on arrangements, where two arrangements are equivalent if one is a translation of the other. We define a *configuration* to be an equivalence class of arrangements. If configuration  $\sigma$  is a rotation of configuration  $\tau$ , we still consider  $\sigma$  and  $\tau$  to be distinct configurations. That is, for the purpose of monitoring the particle system we maintain a global orientation of the particles, even though each individual particle has no sense of global orientation.

We will let capital letters refer to particles and lower case letters refer to locations on the triangular lattice  $\Gamma$ , e.g., “particle  $P$  at location  $\ell$ .” For a particle  $P$ , we use  $N(P)$  to denote the set of particles adjacent to  $P$ , where by *adjacent* we mean connected by a lattice edge. Similarly, for a location  $\ell$ , let  $N(\ell)$  denote the set of particles adjacent to  $\ell$ , excluding any particle at location  $\ell$ . For a particle  $P$ , we will use  $n(P)$  to denote the six locations in the neighborhood of  $P$ , excluding the location  $P$  occupies; for a location  $\ell$ ,  $n(\ell)$  similarly denotes the six locations adjacent to  $\ell$ , excluding  $\ell$  itself. For locations  $\ell$  and  $\ell'$ , by  $N(\ell \cup \ell')$  we mean  $(N(\ell) \cup N(\ell')) \setminus \{\ell, \ell'\}$ ; the same holds for  $n(\ell \cup \ell')$ . In an abuse of notation, we occasionally write  $N(\cdot) \setminus \{\ell\}$ , which removes a location from a set that only contains particles; as one might expect, this notation refers to removing any particle that is at location  $\ell$  from the set of particles  $N(\cdot)$ .

By an *edge* of a configuration  $\sigma$  we mean an edge of  $\Gamma$  such that both incident vertices are occupied by particles. Similarly, by a *triangle* of  $\sigma$  we mean a triangular face of  $\Gamma$  with all three vertices occupied by particles. We denote the number of edges of  $\sigma$  is  $e(\sigma)$  and the number of triangles by  $t(\sigma)$  Throughout, by a *path* or a *cycle* we mean a path or cycle in the underlying graph  $\Gamma$  where all vertices are occupied by particles, and in the case of a cycle, at least one location inside the cycle is unoccupied. Two particles are *connected* if there exists a path between them, and a configuration is *connected* if all pairs of particles are. A *hole* in a configuration is a maximal finite component of adjacent unoccupied locations. A configuration has a hole if and only if it has a cycle. We define the *perimeter*  $p(\sigma)$  of a particle configuration  $\sigma$  to be the sum of the lengths of the walks along all portions of its boundary, including the boundaries of holes.

In general, the configurations we consider are those with all particles contracted. These will be the states of our Markov chain, and in Section 5.4.2 we extensively discuss how to implement our Markov chain for compression in an asynchronous distributed way where particle expansions and contractions may not be consecutive. In particular, when calculating the perimeter of a particle configuration we will ignore any heads of expanded particles and only consider the locations occupied by their tails and by contracted particles. This is for technical reasons that will become clear later.

We specifically focus on connected particle configurations. Because the particles we consider can only communicate with immediate neighbors, if a particle configuration is separated into multiple components there is no way for those components to communicate, interact, or find each other in any reliable way. Furthermore, our current proof techniques require hole-free configurations. Our algorithm doesn't allow new holes to form, and if it begins at a connected configuration with holes we guarantee existing holes will be eliminated.

### 5.2.3 Related Work: Particle Exclusion Processes

As opposed to earlier work in the amoebot model, we use randomization to determine particle movements. When we couple a particle's expansion and contraction steps into one move, the resulting random dynamics are an example of a *particle exclusion process*, where some fixed number of particles on a graph move among vertices, never occupying the same vertex at the same time. There has been a significant body of work analyzing Markov chains that are particle exclusion processes. In fact, the widely-used Comparison Theorem for bounding the mixing time of Markov chains (Theorem 2.6) was first presented in a paper in which it was used to analyze the mixing time of an unbiased exclusion process [43]. There are also close connections between exclusion processes and card shuffling (e.g., [9], and between exclusion processes and statistical physics models (e.g. [111]).

Our setting of the geometric amoebot model requires us to diverge from many common assumptions made about exclusion processes. In order to accomplish desired objectives, the probabilities of particle moves are not fixed ahead of time but are calculated anew in each iteration based on the configuration in the neighborhood of the move. Our random particle dynamics are also constrained to ensure the particle configuration remains connected. This is necessary because the amoebot model limits communication to immediate neighbors, and furthermore on infinite graphs like the one we consider typical exclusion processes will simply result in the diffusion of particles away from each other. These divergences from previous work on exclusion processes are necessary to accomplish our objectives and to be able to translate our Markov chains processes into useful distributed algorithms.

## **5.3 Compression: Overview, Problem Definition, and Preliminaries**

We now focus on the *compression* problem, in which the particle system gathers as tightly together as possible. This phenomenon is often found in natural systems from fire ants [89] to honeybees [19]. While each individual ant or bee cannot view the group as a whole when

soliciting information, it can take cues from its immediate neighbors to achieve cooperation. It is with this motivation that we present a distributed algorithm for compression in the amoebot model of programmable matter.

Specifically, the *compression problem* seeks to reorganize the configuration of a particle system (via movements of particles) such that the system converges to a configuration with small perimeter. We say a particle system is  $\alpha$ -compressed, for  $\alpha > 1$ , if the perimeter of the particle configuration is at most  $\alpha$  times the minimum possible perimeter for those particles.

We present a *Markov chain*  $\mathcal{M}$  for particle compression under the geometric amoebot model that can be directly translated into a *fully distributed, local, asynchronous compression algorithm*  $\mathcal{A}$ , executed by each particle independently without any centralized control. Both  $\mathcal{A}$  and  $\mathcal{M}$  take as input a *bias parameter*  $\lambda$  (where  $\lambda > 1$  favors smaller perimeter) and start from an arbitrary initial configuration for the particles that is connected. In our amoebot model, we assign each configuration  $\sigma$  a Hamiltonian  $H(\sigma) = p(\sigma)$ , the perimeter of the particle configuration;  $\mathcal{M}$  is designed to converge to a distribution that favors small Hamiltonians, proportional to  $e^{-BH(\sigma)} = \lambda^{-p(\sigma)}$ . As  $\lambda$  gets larger (by increasing  $B$ , effectively lowering temperature), we favor configurations with smaller perimeter, causing increasingly compressed configurations.

Markov chain  $\mathcal{M}$  is carefully designed according to the distributed and local nature of the system, so that the particles always stay connected and holes are eliminated and don't reform. Furthermore, we prove  $\mathcal{M}$  is *reversible and ergodic*, meaning many of the standard tools of Markov chain analysis can be applied. While most of these proofs rely only on first principles, we emphasize they are far from trivial; working in a distributed setting necessitates carefully defined protocols for local moves that make proofs challenging.

When the particles execute the local moves of  $\mathcal{M}$  (by each running  $\mathcal{A}$ ) for long enough, the configuration of the particles converges to the stationary distribution of  $\mathcal{M}$ . We prove for all large enough  $\lambda$  there is a constant  $\alpha = \alpha(\lambda) > 1$  such that at stationarity, with all but



exponentially small probability, the particles are  $\alpha$ -compressed, meaning the perimeter of the particle configuration is at most  $\alpha$  times the minimum perimeter (which is  $\Theta(\sqrt{n})$  for systems of  $n$  particles). We additionally show the counterintuitive result that  $\lambda > 1$  is not enough to guarantee compression, even though when  $\lambda > 1$  configurations with smaller perimeters have more weight. In fact, for all  $0 < \lambda < 2.17$ , there is a constant  $\beta < 1$  such that at stationarity with all but exponentially small probability the perimeter is at least a  $\beta$  fraction of the maximum perimeter, which is  $\Theta(n)$  for systems of  $n$  particles. We call such a configuration  $\beta$ -expanded. This implies that for any  $0 < \lambda < 2.17$ , the probability that the particles are  $\alpha$ -compressed is exponentially small for any constant  $\alpha$ .

The key tool used to establish compression is a careful *Peierls argument*, used in statistical physics to study non-uniqueness of limiting Gibbs measures and to determine the presence of phase transitions (see, e.g., [49]), and in computer science to establish slow mixing of Markov chains (see, e.g., [14]).

### 5.3.1 Related Work

Our work on compression was originally inspired in part by the *Ising model* of statistical physics [66], which has been widely studied. In this model, all vertices of some graph are assigned a positive or negative spin, and a *temperature* parameter governs how likely it is for neighboring particles to have the same spin. For certain temperatures, we see *clustering*, where large regions of the graph have the same spin. In our model, the graph we consider is the infinite triangular lattice  $\Gamma$ , we can view locations occupied by particles as having positive spin, unoccupied locations as having negative spin, and our parameter  $\lambda$ , correlated with inverse temperature, governs the likelihood of adjacent positive spins (adjacent particles). Solving the compression problem for particles corresponds to forming a cluster of positive spins in the Ising model with *fixed magnetization*, where the total number of vertices with each spin does not change. Our work diverges from the fixed magnetization Ising model by requiring that particles only move to adjacent locations and

the particle configuration remains connected, constraints not typically considered for Ising models but necessary for distributed implementations in self-organizing particle systems.

In distributed computing, the rendezvous (or gathering) problem seeks to gather mobile agents together on some node of a graph (see, e.g., [5] and the references within). In comparison, our particles follow the exclusion principle, and hence would not be able to gather at a single node, and are computationally simpler than the mobile agents considered.

Nature offers a variety of examples in which gathering and cooperative behavior are apparent. For example, social insects often exhibit compression-like characteristics in their collective behavior: fire ants form floating rafts [89], cockroach larvae perform self-organizing aggregation [68, 101], and honey bees choose hive locations based on a decentralized process of swarming and recruitment [19].

Lastly, in [38, 42], algorithms for hexagon shape formation in the amoebot model were presented. Although a hexagon satisfies our definition of compression, the Markov chain-based algorithm we present takes a fully decentralized and local approach. This is naturally self-stabilizing, forgoing the need for a seed particle that may coordinate or initiate some underlying organization of the set of particles, as required in [38] and even more critically in [42].

### 5.3.2 Formalizing Compression: Perimeter and Edges

Our objective is to find a solution to the particle compression problem. There are many ways to formalize what it means for a particle system to be *compressed*. For example, one could try to minimize the diameter of the system, maximize the number of edges, or maximize the number of triangles. We choose to define compression in terms of minimizing the *perimeter*. We prove that for connected configurations with no holes (the states we eventually reach), minimizing perimeter, maximizing the number of edges, and maximizing the number of triangles are all equivalent and are stronger notions of compression than minimizing the diameter.

Recall we define the *perimeter*  $p(\sigma)$  of a connected particle configuration  $\sigma$  to be the sum of the lengths of the walks along all portions of its boundary, including the boundaries of holes. Furthermore, we ignore the heads of any expanded particles when calculating perimeter; this constraint and the reasons for it are discussed in Section 5.4.2. In an abuse of notation, we use the term *perimeter* to refer both to the total length  $p(\sigma)$  of these walks and the walks themselves. By *outer perimeter* we mean the unique portion of the perimeter that encloses all particles. If a configuration has no holes, then the perimeter and outer perimeter are the same. We assume any walk  $\mathcal{W}$  along the outer perimeter of a configuration is in the clockwise direction. Note an edge may appear twice in a perimeter walk  $\mathcal{W}$ ; in this case, its length is counted twice in  $p(\sigma)$ .

For a connected configuration of  $n$  particles, the perimeter ranges from a maximum value of  $2n - 2$  when the particles are in their least compressed state (a tree with no induced triangles) to some minimum value  $p_{min}(n) = \Theta(\sqrt{n})$  when the particles are in their most compressed state. It is easy to see  $p_{min}(n) \leq 4\sqrt{n}$ , and we now prove any configuration  $\sigma$  of  $n$  particles has  $p(\sigma) \geq \sqrt{n}$ ; this bound is not tight but suffices for our proofs.

**Lemma 5.1.** *A connected configuration with  $n \geq 2$  particles has perimeter at least  $\sqrt{n}$ .*

*Proof.* We argue by induction on  $n$ . A connected particle system with two particles necessarily has perimeter  $2 \geq \sqrt{2}$ , as claimed. Let  $\sigma$  be any particle configuration with  $n$  particles where  $n > 2$ , and suppose the lemma holds for all configurations with less than  $n$  particles.

First, suppose there is a particle  $Q \in \sigma$  not incident to any triangles of  $\sigma$ . This implies  $Q$  has one, two, or three neighbors, none of which are adjacent. If  $Q$  has one neighbor, removing  $Q$  from  $\sigma$  yields a configuration  $\sigma'$  with  $n - 1$  particles and, by induction, perimeter at least  $\sqrt{n - 1}$ . Thus

$$p(\sigma) = p(\sigma') + 2 \geq \sqrt{n - 1} + 2 \geq \sqrt{n}.$$

If  $Q$  has two neighbors, removing  $Q$  from  $\sigma$  produces two connected particle configurations  $\sigma_1$  and  $\sigma_2$ , where  $\sigma_1$  has  $n_1$  particles,  $\sigma_2$  has  $n_2$  particles, and  $n_1 + n_2 = n - 1$ . Thus,

$$p(\sigma) \geq \sqrt{n_1} + \sqrt{n_2} + 4 > \sqrt{n-1} + 4 > \sqrt{n}.$$

Similarly, if  $Q$  has three neighbors its removal produces three particle configurations with  $n_1$ ,  $n_2$ , and  $n_3$  particles, where  $n_1 + n_2 + n_3 = n - 1$ , and we conclude

$$p(\sigma) \geq \sqrt{n_1} + \sqrt{n_2} + \sqrt{n_3} + 6 > \sqrt{n-1} + 6 > \sqrt{n}.$$

Now, suppose every particle in  $\sigma$  is incident to some triangle of  $\sigma$ , implying there are at least  $\lceil n/3 \rceil$  triangles in  $\sigma$ . An equilateral triangle with side length 1 has area  $\sqrt{3}/4$ , so the perimeter of  $\sigma$  encloses an area of at least  $A = \lceil n/3 \rceil \sqrt{3}/4 \geq \sqrt{3}n/12$ . By the isoperimetric inequality, the minimum perimeter way of enclosing this area, without regard to lattice constraints, is with a circle of radius  $r$  and perimeter  $p$ , where

$$r = \sqrt{\frac{A}{\pi}} = \sqrt{\frac{n\sqrt{3}}{12\pi}}, \quad p = 2\pi r = \sqrt{\frac{\pi n}{\sqrt{3}}} > \sqrt{n}.$$

As the perimeter of  $\sigma$  also encloses an area of at least  $\sqrt{3}n/12$ , it is of length at least  $\sqrt{n}$ . □

When  $n$  is clear from context we omit it and refer to  $p_{min} = p_{min}(n)$  and  $p_{max} = p_{max}(n)$ .

We now define what it means for a particle system to be compressed.

**Definition 5.2.** *For any  $\alpha > 1$ , a connected configuration  $\sigma$  with no holes is  $\alpha$ -compressed if  $p(\sigma) \leq \alpha \cdot p_{min}$ .*

We prove in Section 5.5.2 that our algorithm, when executed for a sufficiently long time, achieves  $\alpha$ -compression with all but exponentially small probability for any constant  $\alpha > 1$ , provided  $n$  is sufficiently large. We note  $\alpha$ -compression implies the diameter of

the particle system is also  $O(\sqrt{n})$ , so our definition of  $\alpha$ -compression is stronger than defining compression in terms of diameter.

In order to minimize perimeter using only simple local moves, we exploit the following relationship. Because we will show that our algorithm eventually reaches, and remains, in the set of particle configurations with no holes, our statement and proof only consider that case.

**Lemma 5.3.** *For a connected particle configuration  $\sigma$  with no holes,  $e(\sigma) = 3n - p(\sigma) - 3$ .*

*Proof.* We count particle-edge incidences, of which there are  $2e(\sigma)$ . Counting another way, every particle has six incident edges, except for those on the perimeter. Consider any traversal  $\mathcal{W}$  of the perimeter; at each particle, the exterior angle is 120, 180, 240, 300, or 360 degrees. These correspond to the particle “missing” 1, 2, 3, 4, or 5 of its possible six incident edges, or  $degree/60 - 1$  missing edges. If  $\mathcal{W}$  visits the same particle multiple times, we count the appropriate exterior angle, and the corresponding missing edges, each time  $\mathcal{W}$  visits. From a well-known result about simple polygons with  $p(\sigma)$  sides, the sum of exterior angles along  $\mathcal{W}$  is  $180p(\sigma) + 360$  degrees. Summing up the ‘missing’ edges across all particles on the boundary of  $\sigma$ , we see that there are

$$(180p(\sigma) + 360)/60 - p(\sigma) = 2p(\sigma) + 6$$

total missing edges. This implies there are  $6n - 2p(\sigma) - 6$  total particle-edges incidences, so  $2e(\sigma) = 6n - 2p(\sigma) - 6$ . □

We briefly note that minimizing perimeter is also equivalent to maximizing triangles.

**Lemma 5.4.** *For a connected particle configuration  $\sigma$  with no holes,  $t(\sigma) = 2n - p(\sigma) - 2$ .*

*Proof.* The proof is nearly identical to that of Lemma 5.3, counting particle-triangle incidences instead, of which there are  $3t(\sigma)$ . Counting another way, every particle has six incident triangles, except for those on the perimeter. Consider any traversal  $\mathcal{W}$  of the

perimeter; at each particle, the exterior angle is 120, 180, 240, 300, or 360 degrees. These correspond to the particle “missing” 2, 3, 4, 5, or 6 of its possible six incident triangles, or  $\text{degree}/60$  missing triangles. If  $\mathcal{W}$  visits the same particle multiple times, count the appropriate exterior angle at each visit. The sum of exterior angles along  $\mathcal{W}$  is  $180p(\sigma) + 360$ , so in total particles on the perimeter are missing  $3p(\sigma) + 6$  triangles. This implies there are  $6n - 3p(\sigma) - 6$  particle-triangle incidences, so  $3t(\sigma) = 6n - 3p(\sigma) - 6$ .  $\square$

The above lemmas give the following corollary.

**Corollary 5.5.** *A connected particle configuration  $\sigma$  with no holes and minimum perimeter is also a configuration with the maximum number of edges and the maximum number of triangles.*

Because these three notions of compression are equivalent, for simplicity we state our algorithm in terms of minimizing the number of edges but prove our compression results in terms of perimeter. When we originally presented these results [21], we stated our algorithm in terms of triangles, but do not do so here.

## 5.4 Algorithms for Compression

Our objective is to demonstrate how stochastic algorithms can provably achieve compression, focusing on self-organizing particle systems on the triangular lattice  $\Gamma$ . Our algorithm is based on Markov chain principles that will enable us to prove rigorous results about its behavior. Remarkably, it does not even require the particles to communicate more than one bit of information to each other, even though the amoebot model allows for such exchanges; at any activation, a particle only needs to know which of its neighboring locations are occupied and which, if any, of those neighbors are expanded.

Our algorithm achieves compression by making a particle more likely to move into a position where it has more neighbors, that is, where it forms more edges with neighboring particles. Specifically, a bias parameter  $\lambda$  controls how strongly the particles favor having

neighbors;  $\lambda > 1$  corresponds to favoring more neighbors, while  $\lambda < 1$  corresponds to disfavoring neighbors. As Lemma 5.3 shows, locally favoring more neighbors is equivalent to globally favoring a shorter perimeter; this is the relationship we exploit to obtain particle compression.

#### 5.4.1 The Markov Chain $\mathcal{M}$

We start with two key properties that enable a particle to move from location  $\ell$  to adjacent location  $\ell'$  without disconnecting the particle system or forming a hole. Let  $\mathbb{S} = N(\ell) \cap N(\ell')$  be the set of particles adjacent to both  $\ell$  and  $\ell'$ ; note  $|\mathbb{S}| \in \{0, 1, 2\}$ .

**Property 1.**  $|\mathbb{S}| \in \{1, 2\}$  and every particle in  $N(\ell \cup \ell')$  is connected to a particle in  $\mathbb{S}$  by a path through  $N(\ell \cup \ell')$ .

**Property 2.**  $|\mathbb{S}| = 0$ ,  $\ell$  and  $\ell'$  each have at least one neighbor, all particles in  $N(\ell) \setminus \{\ell'\}$  are connected by paths within this set, and all particles in  $N(\ell') \setminus \{\ell\}$  are connected by paths within this set.

These properties capture precisely the structure required to maintain particle connectivity and prevent new holes from forming. Additionally, both are symmetric for  $\ell$  and  $\ell'$ , necessary for reversibility. However, they are not so restrictive as to limit the movement of particles and prevent compression from occurring. That is, we will see that after a burn-in phase to eliminate any holes that moves satisfying these properties suffice to transform any configuration into any other.

We now define our Markov chain  $\mathcal{M}$  for compression. The state space  $\Omega$  of  $\mathcal{M}$  is the set of all connected configurations of  $n$  contracted particles, and the rules and probabilities given in Algorithm  $\mathcal{M}$  define the transitions between states. Later, in Section 5.4.2, we will show how to view this Markov chain as a local, distributed, asynchronous algorithm  $\mathcal{A}$ . Both  $\mathcal{M}$  and  $\mathcal{A}$  take as input a bias parameter  $\lambda > 1$  and begin at an arbitrary connected starting configuration  $\sigma_0 \in \Omega$ .

---

**Algorithm  $\mathcal{M}$ : Markov Chain for Compression**

---

**Beginning at any connected configuration  $\sigma_0$  of  $n$  contracted particles, repeat:**

- 1: Select particle  $P$  uniformly at random from among all particles; let  $\ell$  be its location. Choose a neighboring location  $\ell'$  and  $q \in (0, 1)$  uniformly at random.
  - 2: **if**  $\ell'$  is unoccupied **then**
  - 3:      $P$  expands to simultaneously occupy  $\ell$  and  $\ell'$ .
  - 4: **else** Return to Step 1.
  - 5: Let  $e = |N(\ell)|$  be the number of neighbors  $P$  had when it was contracted at  $\ell$ , and let  $e' = |N(\ell')|$  be the number of neighbors  $P$  would have if it contracts to  $\ell'$ .
  - 6: **if** (1) location  $\ell$  does not have five neighboring particles, (2) locations  $\ell$  and  $\ell'$  satisfy Property 1 or Property 2, and (3)  $q < \lambda^{e'-e}$  **then**
  - 7:      $P$  contracts to  $\ell'$ .
  - 8: **else**  $P$  contracts back to  $\ell$ .
- 

In Markov chain  $\mathcal{M}$ , note that a constant number of random bits suffice to generate  $q$ , as only a constant precision is required (given that  $e' - e$  is an integer in  $[-3, 3]$  and  $\lambda$  is a constant). In Step 6, Condition (1) ensures no holes form, Condition (2) ensures the particle system stays connected and  $\mathcal{M}$  is reversible, and Condition (3) ensures the particle moves happen with probabilities such that  $\mathcal{M}$  converges to the desired distribution.

#### 5.4.2 The Local Algorithm $\mathcal{A}$

We now present the local, distributed, asynchronous algorithm that each particle runs. Recall from Section 5.2.1 that during a single activation of a particle  $P$ ,  $P$  can perform an arbitrary amount of computation and at most one expansion or contraction. In particular,  $P$  cannot do both an expansion and a contraction in one activation as  $\mathcal{M}$  does in a single state transition. Thus, in  $\mathcal{A}$ , we decouple a single state transition of  $\mathcal{M}$  into two (not necessarily consecutive) particle activations and carefully handle the way in which the particle's neighborhood may change between its two activations.

Each particle  $P$  continuously runs Algorithm  $\mathcal{A}$ , executing Steps 1–7 if  $P$  is contracted, and Steps 8–12 if  $P$  is expanded. Conditions (1)–(3) in Step 10 of  $\mathcal{A}$  are the same as those in Step 6 of  $\mathcal{M}$ , considering the set  $N^*$ , which treats expanded particles as if they are still contracted at their tail location, rather than the set  $N$  of occupied neighboring locations.



---

**Algorithm  $\mathcal{A}$ :** Local, Distributed, Asynchronous Algorithm for Compression Run Independently by each Particle  $P$

---

**If  $P$  is contracted:**

- 1: Let  $\ell$  denote  $P$ 's current location.
- 2: Particle  $P$  chooses neighboring location  $\ell'$  uniformly at random from the six possible choices, and generates a random number  $q \in (0, 1)$ .
- 3: **if**  $\ell'$  is unoccupied and  $P$  has no expanded neighboring particle at  $\ell$  **then**
- 4:      $P$  expands to simultaneously occupy  $\ell$  and  $\ell'$ .
- 5:     **if** there are no expanded particles adjacent to  $\ell$  or  $\ell'$  **then**
- 6:          $P$  sets  $flag = TRUE$  in its local memory.
- 7:     **else**  $P$  sets  $flag = FALSE$ .

**If  $P$  is expanded:**

- 8: Let  $N^*(\cdot) \subseteq N(\cdot)$  be the set of neighboring particles excluding any heads of expanded particles.
  - 9: Let  $e = |N^*(\ell)|$  be the number of neighbors  $P$  had when it was contracted at  $\ell$ , and let  $e' = |N^*(\ell')|$  be the number of neighbors  $P$  would have if it contracts to  $\ell'$ .
  - 10: **if** (1)  $e \neq 5$ , (2) locations  $\ell$  and  $\ell'$  satisfy Property 1 or Property 2 with respect to  $N^*(\cdot)$ , (3)  $q < \lambda^{e'-e}$ , and (4)  $flag = TRUE$  **then**
  - 11:      $P$  contracts to  $\ell'$ .
  - 12: **else**  $P$  contracts back to  $\ell$ .
- 

Note only a constant number of bits are needed to produce  $q$ , as  $\lambda$  is a constant and a particle move changes the number of edges by at most a constant amount. The additional Condition (4) ensures  $P$  is the only particle in its neighborhood potentially moving to a new position since it last expanded. If condition 4 is satisfied, then any expanded particles in  $n(\ell \cup \ell')$  must have expanded after  $P$  did; this means they have  $flag = FALSE$  and will contract to their original location in their next activation, justifying our use of  $N^*$  which ignores heads of adjacent expanded particles.

Any conflicts arising from two particles concurrently attempting to expand into the same location are assumed to be resolved arbitrarily. Hence, any concurrent movements will cover pairwise disjoint neighborhoods and the respective actions will be mutually independent. Following the classical asynchronous model [80], for any starting configuration  $\sigma_0$  and any concurrent execution of  $\mathcal{A}$  that reaches an arbitrary configuration  $\sigma$ , there is a sequence of atomic actions that also reaches  $\sigma$ . So it suffices to consider particle expansions

and contractions as atomic actions instead of being executed in parallel.

We now formally explore the connection between Markov chain  $\mathcal{M}$  and distributed algorithm  $\mathcal{A}$ , which decouples a single transition of  $\mathcal{M}$  into a pair of consecutive expansion and contraction particle actions. Let  $P$  be a particle that eventually moves from location  $\ell$  to  $\ell'$  by expanding to occupy both positions at some time  $t$  and contracting to  $\ell'$  at some time  $t' > t$  according to an execution of  $\mathcal{A}$ . Since  $P$  eventually completes its movement to  $\ell'$ , there must have been no expanded particles adjacent to  $\ell$  or  $\ell'$  at time  $t$  (by Step 6 and Condition (4) of Step 10 in  $\mathcal{A}$ ). Any other particle  $Q$  which expands into the neighborhood of  $P$  in the time interval  $(t, t')$  will see that  $P$  is expanded and set its flag to *FALSE* in Step 7. Recall from Section 5.2.1 that a particle can differentiate between a neighbor's head and tail. By ignoring the heads of any such expanded particle  $Q$  in Steps 8–10,  $P$  can make decisions as if  $Q$  had never moved while  $Q$  eventually contracts back to its original position during its next activation. Thus, the neighborhood of  $P$  remains effectively undisturbed in the interval  $(t, t')$ , allowing  $\mathcal{A}$  to faithfully emulate  $\mathcal{M}$ .

Local algorithm  $\mathcal{A}$  and Markov chain  $\mathcal{M}$  are very closely related. First, consider any sequential execution of atomic actions that leads to configuration  $\sigma'$  in  $\mathcal{A}$ , in which some particles may be expanded and some particles may be contracted. Let configuration  $\sigma$  be obtained from  $\sigma'$  by preserving the locations of all contracted particles and considering every expanded particle to be contracted at its tail (heads of expanded particles are discounted because they are considered ‘exploratory’). Then there exists a sequence of transitions in  $\mathcal{M}$  that reaches  $\sigma$ . The perimeter  $p(\sigma')$  ignores heads of expanded particles, and so  $p(\sigma) = p(\sigma')$ . Conversely, every sequence of transitions in  $\mathcal{M}$  that reaches a configuration  $\sigma$  directly corresponds to a sequence of atomic actions in  $\mathcal{A}$  also leading to  $\sigma' = \sigma$ , where again  $p(\sigma) = p(\sigma')$ . This means that proving  $\alpha$ -compression for  $\sigma$  also implies  $\alpha$ -compression for  $\sigma'$ , and vice-versa. Hence, we can use  $\mathcal{M}$  and respective Markov chain tools and techniques in order to analyze the correctness of algorithm  $\mathcal{A}$ . In particular, because we show  $\alpha$ -compression for  $\mathcal{M}$  for all  $\alpha > 1$ , this also then implies  $\alpha$ -compression

for  $\mathcal{A}$  for all  $\alpha > 1$ . For this reason, in the subsequent sections, we focus on analyzing  $\mathcal{M}$ .

We note that, under the assumptions of the asynchronous model of distributed computing, one cannot typically assume the next particle to be activated is equally likely to be any particle, as we do in Step 1 of  $\mathcal{M}$ . We make this assumption in order to be able to explicitly calculate the stationary distribution of  $\mathcal{M}$  so that we can provide rigorous guarantees about its structure, but do not expect the behavior of the system would be substantially different if this requirement was relaxed.

To justify this random activation assumption, we note that random sequences of particle activations can be approximated using Poisson clocks with mean 1 (the analysis can be modified to accommodate each clock having its own constant mean; however, for ease of presentation, we assume here that they are all identically distributed). That is, each particle activates and executes Algorithm  $\mathcal{A}$  at a random real time drawn from the exponential distribution  $e^{-t}$ . After each action, the particle then computes another random time drawn from the same distribution  $e^{-t}$  and activates again after that amount of time has elapsed. The exponential distribution is unique in that, if particle  $P$  has just activated, it is equally likely that any particle will be the next particle to activate, including particle  $P$  (see, e.g., [55]). Moreover, the particles update without requiring knowledge of any of the other particles' clocks. Similar Poisson clocks are commonly used to describe physical systems that perform updates in parallel in continuous time.

While this shows our Markov chain  $\mathcal{M}$  can be translated into a fully local, distributed algorithm with the same behavior, such an implementation is not always possible in general. Any Markov chain for particle systems that inherently relies on non-local moves of particles or has transition probabilities relying on non-local information cannot be executed by a local, distributed algorithm. Moreover, most distributed algorithms for amoebot systems are not stochastic and thus cannot be meaningfully described as Markov chains; see, e.g., the mostly deterministic algorithms in [40, 41].

### 5.4.3 Obliviousness and Robustness of $\mathcal{M}$ and $\mathcal{A}$

Our algorithm for compression is the first (nearly) *oblivious* algorithm for self-organizing particle systems; that is, whenever a particle is activated, it remembers no information from past activations and decides what to do based only on its observation of the current environment. Obliviousness in mobile robots has been considered in a large number of settings, and a summary of this work can be found in [57]. In practical settings, oblivious robots are desirable because they do not require persistent memory and often are self-stabilizing and fault-tolerant; theoretically, they are of great interest because they are computationally weak at an individual level but can still collectively accomplish sophisticated goals.

Previous works on self-organizing particle systems under the amoebot model (e.g., [42, 36, 35]) make heavy use of particle memory and communication between particles to achieve their tasks. In our Markov chain algorithm for compression, however, if we consider an activation of a particle to be an expansion followed by a contraction (as in  $\mathcal{M}$ ), then our algorithm is oblivious; particles depend only on the current configuration of their neighborhood to make decisions. Even if an activation of a particle is defined to include at most one expansion or contraction (as in  $\mathcal{A}$  and Section 5.2.1), our algorithm is *nearly* oblivious; particles need only store the value of their *flag* variable as one bit of information between their expansion and subsequent contraction activations. Furthermore, communication between particles is also limited to at most this one bit: a particle simply needs to check for flags among neighbors and more complicated message do not need to be relayed.

Moreover, our algorithm for compression is the first for self-organizing particle systems to meaningfully consider fault-tolerance. A distributed algorithm's *fault-tolerance* has to do with its ability to achieve its goals despite possible *crash failures* or *Byzantine failures*. In a crash failure, an agent abruptly ceases functioning and may never be resuscitated. These failures are particularly problematic for systems with a single point of failure, as there is no guarantee the critical agent will remain non-faulty and no guarantee that its memory and role could be assumed by another agent if it crashes. In a Byzantine failure,

some fraction of the agents are malicious and execute arbitrary behavior in an effort to stop the non-faulty portion of the system from achieving its task.

Previous works on self-organizing particle systems have not addressed either type of possible faults, and many of the proposed algorithms are susceptible to complete failure if even a single particle crashes. In contrast, if one or more particles were to crash in our algorithm for compression, they would cease moving and act as fixed points around which the remaining particles would simply continue to compress. For the more adversarial setting of Byzantine failures, we speculate that the malicious particles could affect the overall compression of the system by trying to expand away from where the system is aggregating; however, if the fraction of malicious particles is small, this will not have a large effect. Furthermore, since our algorithm is (nearly) oblivious and does not use particle communication, other than particles checking the flags of neighbors, the malicious particles are unable to “lie” or otherwise try to corrupt healthy particles’ behaviors.

All of these features – obliviousness, limited communication, and robustness – mark a significant step in algorithmic foundations for self-organizing particle systems towards processes that are more practically realizable. Real-world programmable matter systems are prone to faults, which can be amplified by imperfect communication or memory, and our algorithm minimizes opportunities for such errors to occur and propagate.

#### 5.4.4 Invariants for Markov Chain $\mathcal{M}$

Now that we have described and discussed algorithm  $\mathcal{A}$  and shown that it is a distributed implementation of Markov chain  $\mathcal{M}$ , we will perform the rest of our analysis directly on  $\mathcal{M}$ . We begin by showing that  $\mathcal{M}$  maintains certain invariants.

**Lemma 5.6.** *If the particle system is initially connected, during the execution of Markov chain  $\mathcal{M}$  it remains connected.*

*Proof.* Consider one iteration of  $\mathcal{M}$  where a particle  $P$  moves from location  $\ell$  to location  $\ell'$ . Let  $\sigma$  be the configuration before this move, and  $\sigma'$  the configuration after. We show if  $\sigma$  is

connected, then so is  $\sigma'$ .

A move of particle  $P$  from  $\ell$  to  $\ell'$  occurs only if  $\ell$  and  $\ell'$  are adjacent and satisfy Property 1 or Property 2. First, suppose they satisfy Property 1. If  $\sigma$  is connected, then for every particle  $Q$  there exists some path  $\mathcal{P} = (P = P_1, P_2, \dots, P_k = Q)$  from  $P$  to  $Q$  in  $\sigma$ . By Property 1, since  $P_2 \in N(\ell)$ , there exists a path from  $P_2$  to a particle  $S \in \mathbb{S}$  that is entirely contained in  $N(\ell)$ . After  $P$  moves to location  $\ell'$ , it remains connected to particle  $Q$  by a (not necessarily simple) walk that first travels to  $S$ , then travels through  $N(\ell)$  to  $P_2$ , and finally follows  $\mathcal{P}$  to  $Q$ . This implies  $P$  is connected to all particles from location  $\ell'$ , so  $\sigma'$  is connected via paths through  $P$ .

Next, assume locations  $\ell$  and  $\ell'$  satisfy Property 2. Let  $Q, Q' \neq P$  be particles; we show that if  $\sigma$  is connected, then  $Q$  and  $Q'$  must be connected by a path not containing  $P$ . If  $\sigma$  is connected, then  $Q$  and  $Q'$  are connected by some path  $\mathcal{P} = (Q = Q_1, Q_2, \dots, Q_k = Q')$ . If this path doesn't contain  $P$  we are done, so suppose this path contains  $P$ , that is,  $Q_i = P$  for some  $i \in \{2, \dots, k-1\}$ . Both  $Q_{i-1}$  and  $Q_{i+1}$  are neighbors of  $\ell$ , and by Property 2 all neighbors of  $\ell$  are connected by a path in  $N(\ell)$ . Thus  $\mathcal{P}$  can be augmented to form a (not necessarily simple) walk  $\mathcal{W}$  by replacing  $P$  with a path from  $Q_{i-1}$  to  $Q_{i+1}$  in  $N(\ell)$ . As  $P \notin \mathcal{W}$ , this walk connects  $Q$  and  $Q'$  in  $\sigma'$  without going through  $P$ , as desired. Additionally, because  $\ell'$  has at least one neighbor by Property 2,  $P$  remains connected to at least one particle, and via that particle to all other particles in  $\sigma'$ . Thus  $\sigma'$  is connected.  $\square$

**Lemma 5.7.** *Once reaching a connected configuration with no holes, all subsequent configurations of the execution of Markov chain  $\mathcal{M}$  will remain hole-free.*

*Proof.* Consider one iteration of  $\mathcal{M}$  where a particle  $P$  moves from location  $\ell$  to location  $\ell'$ . Let  $\sigma$  be the configuration before this move, and  $\sigma'$  the configuration after. We show if  $\sigma$  is hole-free, then so is  $\sigma'$ . Recall we assume, by definition, that a *cycle* in  $\sigma$  encircles at least one unoccupied location; a configuration has a hole if and only if it has a cycle. Throughout this proof, we will argue about the existence of cycles rather than the existence of holes.

We first show that any cycle introduced in  $\sigma'$  must contain  $P$ . Suppose, for the sake

of contradiction, this is not the case and  $\sigma'$  has a cycle  $\mathcal{C}$  with  $P \notin \mathcal{C}$ . If  $P$  is removed from location  $\ell'$ , then cycle  $\mathcal{C}$  still exists in  $\sigma' \setminus P$ . If  $P$  is then placed at  $\ell$ , yielding  $\sigma$ , then  $\mathcal{C}$  still exists unless it had enclosed exactly one unoccupied location,  $\ell$ . However, this is not possible as any cycle in  $\sigma' - P$  encircling  $\ell$  would also necessarily encircle neighboring unoccupied location  $\ell'$ . This implies cycle  $\mathcal{C}$  exists in cycle-free configuration  $\sigma$ , a contradiction. We conclude every cycle in  $\sigma'$  must contain  $P$ .

Because particle  $P$  moved from location  $\ell$  to location  $\ell'$  in a valid step of Markov chain  $\mathcal{M}$ , it must be true (by the conditions checked in Step 6 of  $\mathcal{M}$ ) that  $\ell$  has fewer than five neighbors and locations  $\ell$  and  $\ell'$  satisfy Property 1 or Property 2. First, suppose they satisfy Property 2. While  $P$  might momentarily create a cycle when it expands to occupy both locations  $\ell$  and  $\ell'$ , it will then contract to location  $\ell'$ . Suppose  $P$  is part of some cycle  $\mathcal{C} = (P = P_1, P_2, \dots, P_{k-1}, P_k = P)$  in  $\sigma'$ . By Property 2,  $P_2$  and  $P_{k-1}$  are connected by a path in  $N(\ell')$  that doesn't contain  $P$ . Replacing path  $(P_{k-1}, P, P_2)$  in cycle  $\mathcal{C}$  by this path in  $N(\ell')$  yields a (not necessarily simple) cycle  $\mathcal{C}'$  in  $\sigma'$  not containing  $P$ , a contradiction.

Next, suppose  $\ell$  and  $\ell'$  satisfy Property 1. Because particle  $P$  moved from  $\ell$  to  $\ell'$  in a valid step of  $\mathcal{M}$ , location  $\ell$  must at most four neighbors in  $\sigma$ . This means that in  $\sigma'$ , location  $\ell$  has at most 5 neighbors – its original neighbors plus  $P$  at location  $\ell'$  – and thus is adjacent to at least one unoccupied location. Suppose there exists some cycle  $\mathcal{C} = (P = P_1, P_2, \dots, P_{k-1}, P_k = P)$  in  $\sigma'$ . This cycle encircles at least one occupied location, and in fact it encircles at least one unoccupied location  $\ell'' \neq \ell$ : since  $\ell$  is adjacent to another unoccupied location in  $\sigma'$ , it cannot be the case that  $\ell$  is the only unoccupied location inside  $\mathcal{C}$ . If there exists a path between  $P_2$  and  $P_{k-1}$  in  $N(\ell')$ , the argument from the first case applies and we are done. Otherwise, without loss of generality, it must be that  $|\mathbb{S}| = 2$  and there exist paths in  $N(\ell \cup \ell')$  from  $P_{k-1}$  to  $S_1 \in \mathbb{S}$  and from  $P_2$  to  $S_2 \in \mathbb{S}$ , with  $S_1 \neq S_2$ . There then exists a (not necessarily simple) cycle  $\mathcal{C}^*$  in  $\sigma$  obtained from  $\mathcal{C}$  by replacing path  $(P_{k-1}, P, P_2)$ , where  $P$  is in location  $\ell'$ , with path  $(P_{k-1}, \dots, S_1, P, S_2, \dots, P_2)$ , where  $P$  is in location  $\ell$ .  $\mathcal{C}^*$  is a valid cycle in  $\sigma$  because it

encircles unoccupied location  $\ell'' \neq \ell$ . This is a contradiction because  $\sigma$  has no cycles. We conclude that, in all cases,  $\sigma'$  has no cycles, and thus has no holes.  $\square$

#### 5.4.5 Eventual Ergodicity of Markov chain $\mathcal{M}$

The state space  $\Omega$  of our Markov chain  $\mathcal{M}$  is the set of all connected configurations of  $n$  contracted particles, and Lemma 5.6 ensures that we always stay within this state space. The initial configuration  $\sigma_0$  of  $\mathcal{M}$  may or may not have holes. By Lemma 5.7, once a hole-free configuration is reached,  $\mathcal{M}$  remains in the part of the state space consisting of all hole-free connected configurations, which we call  $\Omega^*$ . In this section, we prove that from any starting state  $\mathcal{M}$  always reaches  $\Omega^*$ . Furthermore, we prove that  $\Omega^*$  is connected, that is, once  $\mathcal{M}$  reaches  $\Omega^*$ ,  $\mathcal{M}$  is irreducible on this smaller state space. As  $\mathcal{M}$  is also aperiodic, we can conclude it is eventually ergodic on  $\Omega^*$ , a necessary precondition for all of the Markov chain analysis to follow.

We emphasize the details of these proofs are far from trivial. However, they have been substantially simplified and clarified from the originally published conference version of these results [21], where the proof of ergodicity required over 10 single-spaced pages of detailed analysis. Figure 5.2 illustrates one difficulty. It depicts a hole-free particle configuration for which there exist no valid moves satisfying Property 1; the only valid moves satisfy Property 2. Thus if moves satisfying Property 2 are not included, neither  $\Omega$  nor  $\Omega^*$  is connected.

At a high level, we prove that for any configuration  $\sigma$  there exists a sequence of valid particle moves transforming  $\sigma$  into a straight line. Since a straight line is hole-free, this shows that from any initial configuration in  $\Omega$ , there exists a sequence of moves with non-zero probability reaching  $\Omega^*$ , as desired. We then prove  $\mathcal{M}$  is reversible on  $\Omega^*$ . This implies that for any  $\tau \in \Omega^*$  there exists a sequence of valid particle moves transforming a straight line into  $\tau$ . Altogether, this shows for any  $\sigma, \tau \in \Omega^*$  there exists a sequence of valid moves (within  $\Omega^*$ ) transforming  $\sigma$  into  $\tau$ , as desired.






Figure 5.2: A particle configuration for which all valid moves of Markov chain  $\mathcal{M}$  satisfy Property 2; no particle has a valid move satisfying Property 1. This demonstrates the subtlety of the Markov chain rules we have defined.

We will let  $m_1$  be the vertical lattice line containing the leftmost particle(s) in  $\sigma$ . We label the subsequent vertical lattice lines as  $m_2, m_3, m_4$ , and so on. The process for moving the particles into one straight line is a sweep line algorithm, an approach often used in computational geometry [58, 105]. We first consider the particles in leftmost vertical line  $m_1$ , then the particles in  $m_2$ , and so on. When considering line  $m_i$ , we maintain the following invariants:

**Invariants:**

1. All particles left of  $m_i$  form lines stretching down and left.
2. Each such line stretches down and left from a particle in  $m_i$  that has an empty location directly below it.

Figure 5.3a gives an example of a particle configuration and a line  $m_i$  satisfying these invariants. In the remainder of this subsection, we describe how to, starting in a configuration in which the invariants are satisfied for  $m_i$ , find a sequence of valid particle moves after which  $m_{i+1}$  satisfies the invariants. For the configuration in Figure 5.3a, the configuration obtained after first ensuring  $m_{i+1}$  satisfies Invariant 1 is shown in Figure 5.3b, and the configuration after ensuring  $m_{i+1}$  also satisfies Invariant 2 is shown in Figure 5.3c.

Throughout this subsection, a *component of line  $m_i$*  will refer to a maximal collection of particles in  $m_i$  that are connected via paths in  $m_i$ . For example, in Figure 5.3a,  $m_i$  has four components (from top to bottom: of one, two, three, and one particles, respectively).

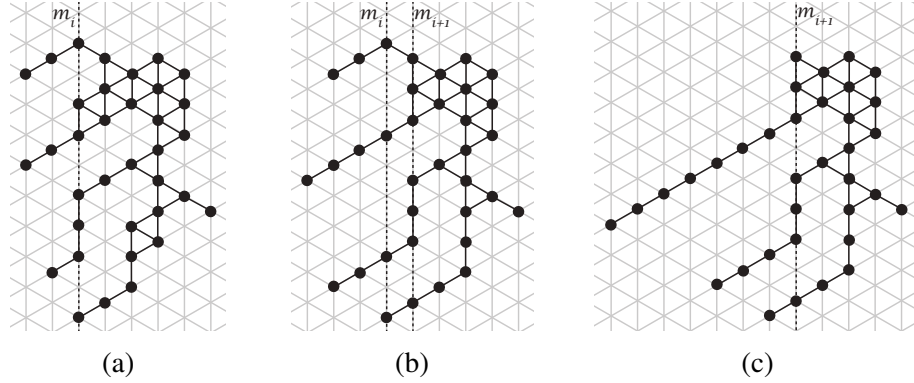


Figure 5.3: (a) An example of a particle configuration and a line  $m_i$  that satisfies both invariants. (b) After a sequence of moves described in Lemma 5.9,  $m_{i+1}$  satisfies Invariant 1. (c) After a sequence of moves described in Lemma 5.10,  $m_{i+1}$  also satisfies Invariant 2.

We begin with a lemma about particle movements that will play a key role.

**Lemma 5.8.** *Suppose particle  $P$  has exactly two neighbors,  $Q_1$  below it and  $Q_2$  above-right of it, and let  $\ell$  be the unoccupied location below-right of  $P$ . There exists a sequence of valid moves, occurring strictly below and right of  $P$ , after which either it is now valid for  $P$  to move to  $\ell$  or some other particle has already moved to  $\ell$ .*

*Proof.* We induct on the number of particles strictly below and right of  $P$ . If there are no such particles, then it is valid (satisfying Property 1) for  $P$  to move from its current location  $\ell_0$  to  $\ell$ . This is because  $N(\ell_0) \cap N(\ell) = \{Q_1, Q_2\}$ , and either these are the only two particles in  $N(\ell_0 \cup \ell)$  (Figure 5.4a) or there is exactly one other particle in  $N(\ell_0 \cup \ell)$  and it is adjacent to  $Q_2$  (Figure 5.4b). Thus the conclusions of the lemma are already satisfied with an empty set of moves.

Suppose there are  $k > 0$  particles strictly below and right of  $P$ , and for all  $0 \leq k' < k$  the lemma holds. If it is already valid for  $P$  to move to  $\ell$ , we are done; an example is given in Figure 5.4c. Otherwise, since  $P$  has fewer than five neighbors, it must be that neither Property 1 nor Property 2 is satisfied. Note  $\mathbb{S} = N(P) \cap N(\ell)$  contains two particles,  $Q_1$  and  $Q_2$ . Because Property 1 doesn't hold, and  $N(P)$  doesn't contain any particles other than those of  $\mathbb{S}$ , it must be that there is a particle  $P'$  in  $N(\ell)$  that is not connected to a particle in  $\mathbb{S}$  by a path within  $N(\ell)$ . Then  $P'$  must occupy the location below-right of  $\ell$ ,

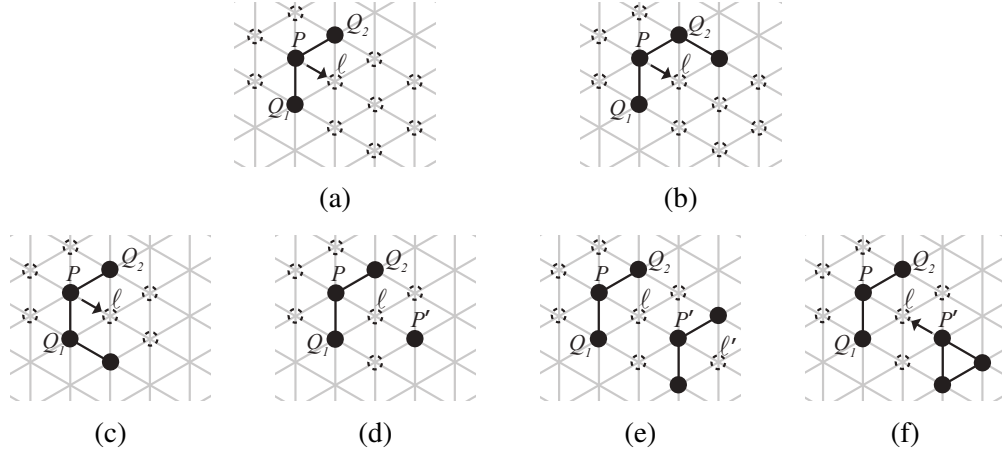


Figure 5.4: Particle positions from the base case (top row) and inductive step (bottom row) of the proof of Lemma 5.8. Particles are represented by black circles, and unoccupied locations are represented by dashed circles. Neighboring particles have a black line drawn between them.

and both locations in  $n(\ell) \cap n(P')$  must be unoccupied; see Figure 5.4d. We now consider  $N(P')$ , which is of size at least one and at most three.

First, we suppose  $N(P')$  is not connected; see Figure 5.4e. In this case,  $P'$  must have exactly two neighbors, one below  $P'$  and the other above-right of  $P'$ , while location  $\ell'$  below-right of  $P'$  is unoccupied. There are fewer than  $k$  particles below and right of  $P'$  because this is a proper subset of the  $k$  particles below and right of  $P$ . By the induction hypothesis, we conclude there is a sequence of moves occurring entirely below and right of  $P'$  after which either it is valid for  $P'$  to move to  $\ell'$  or another particle has moved to  $\ell'$ . In the first case, we let  $P'$  move to  $\ell'$  and afterwards it is valid (satisfying Property 1) for  $P$  to move to  $\ell$ , because  $N(\ell)$  now contains only  $Q_1$  and  $Q_2$ . In the second case, a particle has moved to  $\ell'$  but  $N(P')$  otherwise remains unchanged, causing  $N(P')$  to now be connected, the case we consider next.

Suppose  $N(P')$ , which is of size at least one and at most three, is connected; see Figure 5.4f. Note the current location of  $P'$  and location  $\ell$  satisfy Property 2, so particle  $P'$  can move to  $\ell$ . As  $P'$  and  $\ell$  are below and right of  $P$ , this move satisfies the conclusions of the lemma. □




Figure 5.5: If  $P$  is the topmost particle in a component of  $m_i$  of size at least 2 and its neighborhood is connected, then (a)–(c) are the three possibilities for  $N(P)$ . In all three of these cases, moving  $P$  down-left satisfies Property 1. (d) and (e) show the two cases for subsequently moving  $P$  to a new position such that the invariants still hold for  $m_i$ .

If  $m_i$  satisfies the invariants, we want to give a sequence of moves after which  $m_{i+1}$  also satisfies the invariants. The following lemma will be used towards that goal.

**Lemma 5.9.** *If  $m_i$  satisfies both invariants and has a component of size at least two, there exists a sequence of valid moves that decreases the number of particles in  $m_i$  after which  $m_i$  still satisfies the invariants.*

*Proof.* Consider any component of  $m_i$  of size at least two, and let  $P$  be the topmost particle in this component.  $P$  has a particle below it, no particle above it, and by Invariants 1 and 2 has no particle above-left or below-left of it. The two locations right of  $P$  may or may not be occupied. We consider two cases: when  $N(P)$  is connected, and when it is not.

When  $N(P)$  is disconnected, we invoke Lemma 5.8. It must be that  $P$  has two neighbors that satisfy the conditions of the lemma, and so there exists a sequence of valid moves after which either location  $\ell$  below-right of  $P$  is occupied by another particle or it is valid for  $P$  to move to  $\ell$ . All moves in this sequence occur right of  $P$ , and thus don't affect the invariants for  $m_i$ . If it is now valid for  $P$  to move to  $\ell$ , we make this move and the number of particles in  $m_i$  has decreased, as desired. If another particle has moved to  $\ell$ , then  $N(P)$  is now connected, the next case we consider.

When  $N(P)$  is connected, it must look as in Figure 5.5a, 5.5b, or 5.5c. In all cases, particle  $P$  moving down-left is a valid move that decreases the number of particles in  $m_i$ .

However, Invariant 1 no longer holds for  $m_i$  after this move, so we continue to move particle  $P$  down until it is adjacent to the bottom particle  $Q$  in this component of particles in  $m_i$ . If there is not already a line stretching down and left from  $Q$ , then move  $P$  down once more to start such a line (Figure 5.5d), which is valid because of the invariants for  $m_i$ . If this line stretching down and left from  $Q$  already exists, we note the locations at distances one and two above this line must all be unoccupied. This follows from Invariants 1 and 2 for  $m_i$ : all particles left of  $m_i$  must extend down and left from the bottom particle of some component in  $m_i$ , and the first such particle above  $Q$  is at least two units above  $P$ 's original location and thus at least three units above  $Q$ . Thus, it is valid (satisfying Property 1) to move  $P$  along this line and add it to the end (Figure 5.5e). In all cases, the number of particles in  $m_i$  decreases while the invariants for  $m_i$  remain satisfied, as desired.  $\square$

Lemma 5.9 can be applied iteratively until all components of  $m_i$  are of size one, and all particles left of  $m_i$  form lines stretching down-left from these components of size one. Thus, all particles left of  $m_{i+1}$  form lines stretching down-left, satisfying Invariant 1 for  $m_{i+1}$ . We now consider how to also satisfy Invariant 2 for  $m_{i+1}$ .

**Lemma 5.10.** *If  $m_i$  satisfies both invariants and  $m_{i+1}$  satisfies Invariant 1, then there exists a sequence of valid moves after which  $m_{i+1}$  satisfies both invariants.*

*Proof.* Because our particle configuration is connected, each line left of  $m_{i+1}$  is connected to some particle in  $m_{i+1}$ . However, the line may not stretch down and left from this particle or this particle may not have an empty location below it, as is required by Invariant 2. Consider any component of  $m_{i+1}$  which is adjacent to at least one line left of  $m_{i+1}$  stretching down-left. To satisfy Invariant 2, we merge all such lines into one, stretching down-left from the bottom particle  $Q$  in this component. First, we move the lowest line so that it is stretching down-left from  $Q$ . An entire line can be moved down one unit by first moving the rightmost particle in this line (the particle in line  $m_i$ ) down one unit, and then by subsequently moving the remaining particles down one unit from right to left (for an example of

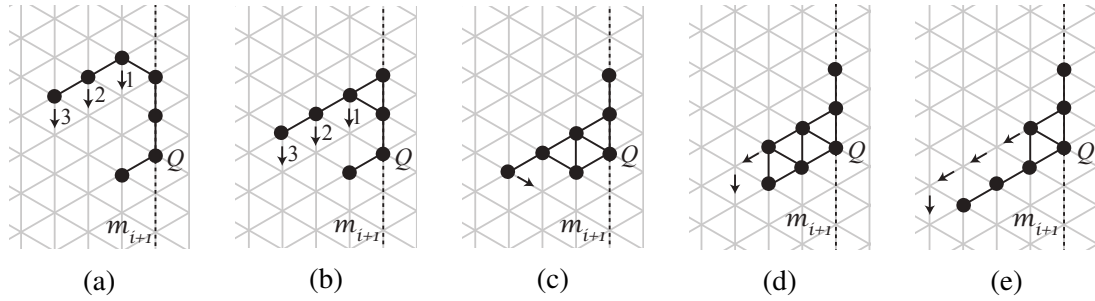


Figure 5.6: The process of merging two lines stretching down-left from the same component of  $m_{i+1}$  in order to satisfy Invariant 2. In (a) and (b), the moves must occur in the order listed.

this downward movement of a line, see Figure 5.6a). This can be repeated until this lowest line is in the desired position, stretching down and left from  $Q$ .

Iteratively consider the next lowest line. As before, we move this line down one unit at a time by moving the particles each down once from right to left until the line is flush with the bottommost line (Figures 5.6a–5.6c). The particles in this line can then easily be added to the bottommost line one at a time, from left to right, as in Figures 5.6c–5.6e. We repeat this line merging process until all particles stretching down-left from this component of  $m_i$  have been reorganized into one line stretching down-left from  $Q$ . After repeating this process for all components in  $m_{i+1}$ , Invariant 2 is satisfied for  $m_{i+1}$ . Invariant 1 still holds for  $m_{i+1}$ , so  $m_{i+1}$  now satisfies both invariants, as claimed.  $\square$

We now combine the previous two lemmas to get the main inductive step for our sweep-line procedure.

**Lemma 5.11.** *If  $m_i$  satisfies both invariants, then there exists a sequence of valid particle moves after which  $m_{i+1}$  also satisfies both invariants.*

*Proof.* Suppose  $m_i$  satisfies both invariants. If there are connected components of two or more particles contained in  $m_i$ , we can iteratively apply Lemma 5.9 to reduce the number of particles in  $m_i$  without affecting the invariants. After this, all components of  $m_i$  consist of one particle. Now all particles left of  $m_{i+1}$  are in lines (possible consisting of just one particle) stretching down-left, satisfying Invariant 1. Next, we can apply Lemma 5.10 to

ensure that  $m_{i+1}$  also satisfies Invariant 2, merging any lines stretching down-left from the same component of  $m_{i+1}$ . Thus, there exists a sequence of valid moves after which  $m_{i+1}$  satisfies both invariants, as claimed.  $\square$

**Lemma 5.12.** *There exists a valid sequence of moves transforming any configuration into a line.*

*Proof.* Initially,  $m_1$  trivially satisfies the invariants because there are no particles left of  $m_1$ . Repeatedly using Lemma 5.11, we obtain a sequence of moves after which the invariants hold for some line  $m_k$  to the right of which there are no particles.

We next note that all particles in  $m_k$  must be in a single component. If this was not the case, then the configuration would not be connected: particles left of  $m_k$  only form lines that are insufficient to connect multiple components of  $m_k$ , and there are no particles right of  $m_k$ . We know that the particle configuration must be connected because we have only made valid particle moves (Lemma 5.6), so this is a contradiction, and  $m_k$  must have a single component.

We repeatedly apply Lemma 5.9 until there is only one particle left in  $m_k$  and line  $m_k$  still satisfies the invariants. At this point the particles form a single line stretching down-left from the single particle in  $m_k$ , and we have given a sequence of valid moves transforming an arbitrary configuration into a line.  $\square$

In particular, this shows that for any configuration there exists a valid sequence of moves transforming it into a configuration with no holes.

**Lemma 5.13.** *Eventually  $\mathcal{M}$  reaches a configuration with no holes, after which no holes are ever introduced again.*

*Proof.* Let  $\sigma_0 \in \Omega$  be the initial (connected) particle configuration given as input to Markov chain  $\mathcal{M}$ . By Lemma 5.12, there is positive probability that  $\mathcal{M}$  will reach  $\Omega^* \subset \Omega$ , the set of hole-free connected particle configurations. Because Lemma 5.12 holds for any configuration, this is also true of each subsequent state  $\sigma_i$ . Since  $\Omega$  is finite,  $\mathcal{M}$  must

eventually reach  $\Omega^*$ , as desired. Finally, by Lemma 5.7, the particle system will remain hole-free for the remainder of  $\mathcal{M}$ 's execution.  $\square$

We present one more lemma before proving  $\mathcal{M}$  is irreducible on  $\Omega^*$  once it reaches  $\Omega^*$ . Recall  $P(\sigma, \tau)$  is the probability of moving from state  $\sigma$  to state  $\tau$  in one step of  $\mathcal{M}$ .

**Lemma 5.14.** *Once  $\mathcal{M}$  reaches  $\Omega^*$ , it is reversible on  $\Omega^*$ . That is, for any two configurations  $\sigma, \tau \in \Omega^*$ , if  $P(\sigma, \tau) > 0$  then  $P(\tau, \sigma) > 0$ .*

*Proof.* Let  $\sigma, \tau \in \Omega^*$  be any two configurations such that  $P(\sigma, \tau) > 0$ . Then  $\sigma$  and  $\tau$  differ by one particle  $P$  that is at location  $\ell$  in  $\sigma$  and at adjacent location  $\ell'$  in  $\tau$ .

Note in  $\tau$ , particle  $P$  at location  $\ell'$  has at most four neighbors. It cannot have six neighbors because location  $\ell$ , which was previously occupied by  $P$  in  $\sigma$ , is now unoccupied. It cannot have five neighbors because otherwise  $\ell'$  would have been a hole in  $\sigma$  when  $P$  was at  $\ell$ , a contradiction to our assumption that  $\sigma \in \Omega^*$ . Because  $P(\sigma, \tau) > 0$ , Property 1 or Property 2 must hold for  $\ell$  and  $\ell'$ . Both properties are symmetric with regard to the role played by  $\ell$  and  $\ell'$ . Thus, if Markov chain  $\mathcal{M}$ , in state  $\tau$ , selects particle  $P$ , location  $\ell \in n(P)$ , and sufficiently small probability  $q$  in Step 1, then because Conditions (1)–(3) of Step 6 are satisfied, particle  $P$  moves to location  $\ell$ . This proves  $P(\tau, \sigma) > 0$ .  $\square$

**Lemma 5.15.** *Once Markov chain  $\mathcal{M}$  reaches  $\Omega^*$ , it connects  $\Omega^*$ , the state space of all connected configurations without holes.*

*Proof.* Let  $\sigma$  and  $\tau$  be any two connected configurations of  $n$  particles with no holes. By Lemma 5.12, there exists a sequence of valid moves transforming  $\sigma$  into a line. By Lemmas 5.12 and 5.14, there exists a sequence of valid moves transforming this line into  $\tau$ .  $\square$

**Corollary 5.16.** *Once  $\mathcal{M}$  reaches  $\Omega^*$ , it is ergodic on  $\Omega^*$ .*

*Proof.* By Lemma 5.15,  $\mathcal{M}$  is irreducible on  $\Omega^*$ . As long as  $n > 1$  then every particle has at least one neighbor, so  $\mathcal{M}$  is aperiodic because at each iteration there is a probability of at least  $1/6$  that a particle proposes moving into an occupied neighboring location so no move is made. Thus, once  $\mathcal{M}$  reaches  $\Omega^*$ , it is ergodic on  $\Omega^*$ .  $\square$



We note that  $\mathcal{M}$  is not irreducible on  $\Omega$ , and thus not ergodic on  $\Omega$ , because it is not possible to get from a hole-free configuration to a configuration with a hole. Ergodicity is necessary to apply tools from Markov chain analysis, as we do in the next subsection.

#### 5.4.6 The Stationary Distribution $\pi$ of $\mathcal{M}$

In this section we determine the stationary distribution of  $\mathcal{M}$ .

**Lemma 5.17.** *If  $\pi$  is a stationary distribution of  $\mathcal{M}$ , then for any  $\sigma \in \Omega \setminus \Omega^*$ ,  $\pi(\sigma) = 0$ .*

*Proof.* For any configuration  $\sigma \in \Omega \setminus \Omega^*$ , there is a positive probability of moving into  $\Omega^*$  in some later time step (Lemma 5.13). For any configuration  $\tau \in \Omega^*$ , there is zero probability of reaching a configuration with holes (Lemma 5.7). If a stationary distribution  $\pi$  were to put any probability mass on states in  $\Omega \setminus \Omega^*$ , over time the total probability mass within  $\Omega \setminus \Omega^*$  would decrease as it leaks into  $\Omega^*$  with no possibility of returning. Thus such a distribution could not be stationary, a contradiction. We conclude that any stationary distribution  $\pi$  has  $\pi(\sigma) = 0$  for all  $\sigma \in \Omega \setminus \Omega^*$ , as claimed.  $\square$

**Lemma 5.18.**  *$\mathcal{M}$  has a unique stationary distribution  $\pi$  given by*

$$\pi(\sigma) = \begin{cases} \frac{\lambda^{e(\sigma)}}{Z} & \sigma \in \Omega^* \\ 0 & \sigma \in \Omega \setminus \Omega^* \end{cases}$$

where  $Z = \sum_{\sigma \in \Omega^*} \lambda^{e(\sigma)}$  is the normalizing constant, also called the partition function.

*Proof.* Lemma 5.17 guarantees that any stationary distribution of  $\mathcal{M}$  has  $\pi(\sigma) = 0$  for configuration  $\sigma \notin \Omega^*$ . Once  $\mathcal{M}$  reaches  $\Omega^*$  (which it is guaranteed to by Lemma 5.13), it is ergodic on  $\Omega^*$  (Lemma 5.16). We conclude, because  $\Omega^*$  is finite, that  $\mathcal{M}$  on  $\Omega^*$  has a unique stationary distribution, and thus  $\mathcal{M}$  on  $\Omega$  also has a unique stationary distribution.

We confirm that  $\pi$  as stated above is this unique stationary distribution by detailed balance. Let  $\sigma$  and  $\tau$  be configurations in  $\Omega^*$  with  $\sigma \neq \tau$  such that  $P(\sigma, \tau) > 0$ . By

Lemma 5.14, also  $P(\tau, \sigma) > 0$ . Suppose particle  $P$  moves from location  $\ell$  in  $\sigma$  to neighboring location  $\ell'$  in  $\tau$ . Let  $e$  be the number of edges formed by  $P$  has when it is in location  $\ell$ , and let  $e'$  be that number when  $P$  is in location  $\ell'$ . This implies  $e(\sigma) - e(\tau) = e - e'$ . If  $\lambda^{e'} \leq \lambda^e$ , then we see that

$$P(\sigma, \tau) = \frac{1}{n} \cdot \frac{1}{6} \cdot \lambda^{e'-e} \quad \text{and} \quad P(\tau, \sigma) = \frac{1}{n} \cdot \frac{1}{6} \cdot 1.$$

In this case we can verify that  $\sigma$  and  $\tau$  satisfy the detailed balance condition:

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{e(\sigma)}}{Z} \cdot \frac{\lambda^{e'-e}}{6n} = \frac{\lambda^{e(\tau)}}{Z \cdot 6n} = \pi(\tau)P(\tau, \sigma).$$

If  $\lambda^{e'} > \lambda^e$ , we can similarly calculate these probabilities to verify detailed balance:

$$P(\sigma, \tau) = \frac{1}{n} \cdot \frac{1}{6} \cdot 1 \quad \text{and} \quad P(\tau, \sigma) = \frac{1}{n} \cdot \frac{1}{6} \cdot \lambda^{e-e'},$$

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{e(\sigma)}}{Z \cdot 6n} = \frac{\lambda^{e(\tau)} \lambda^{e-e'}}{Z \cdot 6n} = \pi(\tau)P(\tau, \sigma).$$

Since the detailed balance condition is satisfied for all  $\sigma, \tau \in \Omega^*$ , it only remains to verify that  $\pi$  is in fact a probability distribution:

$$\sum_{\sigma \in \Omega} \pi(\sigma) = \sum_{\sigma \in \Omega^*} \frac{\lambda^{e(\sigma)}}{Z} + \sum_{\sigma \in \Omega \setminus \Omega^*} 0 = \frac{\sum_{\sigma \in \Omega^*} \lambda^{e(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{e(\sigma)}} = 1.$$

We conclude  $\pi$  is the unique stationary distribution of  $\mathcal{M}$ . □

While it is natural to assume maximizing the number of edges in a particle configuration results in more compression, here we formalize this. We prove  $\pi$  can also be expressed in terms of perimeter.

**Corollary 5.19.** *The stationary distribution  $\pi$  of  $\mathcal{M}$  is also given by*

$$\pi(\sigma) = \begin{cases} \frac{\lambda^{-p(\sigma)}}{Z} & \sigma \in \Omega^* \\ 0 & \sigma \in \Omega \setminus \Omega^* \end{cases}$$

where  $Z = \sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)}$  is the normalizing constant, also called the partition function.

*Proof.* This expression is clearly equal to  $\mathcal{M}$ 's unique stationary distribution when  $\sigma \notin \Omega^*$ , so it only remains to verify the case  $\sigma \in \Omega^*$ . We use Lemma 5.3 and Lemma 5.18:

$$\pi(\sigma) = \frac{\lambda^{e(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{e(\sigma)}} = \frac{\lambda^{3n-p(\sigma)-3}}{\sum_{\sigma \in \Omega^*} \lambda^{3n-p(\sigma)-3}} = \frac{\lambda^{3n-3}}{\lambda^{3n-3}} \cdot \frac{\lambda^{-p(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)}} = \frac{\lambda^{-p(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)}}.$$

□

The the original publication of these results [21] also expressed the stationary distribution in terms of the number of triangles in a configuration, where a *triangle* is a face of  $\Gamma$  that has all three of its vertices occupied by particles, and  $t(\sigma)$  is the number of triangles in configuration  $\sigma$ . We include the following corollary for completeness, but will not use it in subsequent sections.

**Corollary 5.20.** *The stationary distribution  $\pi$  of  $\mathcal{M}$  is also given by*

$$\pi(\sigma) = \begin{cases} \frac{\lambda^{t(\sigma)}}{Z} & \sigma \in \Omega^* \\ 0 & \sigma \in \Omega \setminus \Omega^* \end{cases}$$

where  $Z = \sum_{\sigma \in \Omega^*} \lambda^{t(\sigma)}$  is the normalizing constant, also called the partition function.

*Proof.* This follows from Lemma 5.4 and Corollary 5.19:

$$\pi(\sigma) = \frac{\lambda^{-p(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)}} = \frac{\lambda^{-(2n-t(\sigma)-2)}}{\sum_{\sigma \in \Omega^*} \lambda^{-(2n-t(\sigma)-2)}} = \frac{\lambda^{-2n+2}}{\lambda^{-2n+2}} \cdot \frac{\lambda^{t(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{t(\sigma)}} = \frac{\lambda^{t(\sigma)}}{\sum_{\sigma \in \Omega^*} \lambda^{t(\sigma)}}.$$

□

### 5.4.7 Convergence Time of Markov Chain $\mathcal{M}$

We prove in Section 5.5.2 that when  $\lambda > 2 + \sqrt{2}$ , if Markov chain  $\mathcal{M}$  has converged to its stationary distribution, then with all but exponentially small probability the particle system will be compressed. However, we do not give explicit bounds on the time required for this to occur. We give some experimental results for simulation runs beginning at hole-free configurations in Section 5.5, but believe proving rigorous bounds will be challenging.

A common measure of convergence time of a Markov chain is its *mixing time*, the number of iterations until the distribution is within total variation distance  $\varepsilon$  of the stationary distribution, starting from the worst initial configuration. Getting a polynomial bound on the mixing time of  $\mathcal{M}$ , which achieves compression using local moves in a lattice-based system, is likely to be challenging because of its similarity to physical systems such as the Ising and Potts models, common models of ferromagnetism from statistical physics. For many of these models, Markov chain algorithms that perform local update steps are known to have exponential mixing time, precisely because of a type of compression or clustering that occurs (see, e.g., [110]).

However, mixing time most likely is not the correct measure of our algorithm's convergence. While we prove in later sections that compression occurs after  $\mathcal{M}$  has reached its stationary distribution, we expect (based on simulations and intuition) that compression actually occurs much earlier. Thus, even if it takes exponential time for  $\mathcal{M}$  to converge to its stationary distribution, which is certainly plausible, it may be true that the particles achieve compression after only a polynomial number of steps. When starting from a line of  $n$  particles, our simulations (presented in the next section) indicate that doubling the number of particles consistently results in about a ten-fold increase in iterations until compression is achieved. Based on this, we conjecture the number of iterations until compression occurs is  $\Omega(n^3)$  and  $O(n^4)$ . Furthermore, we do not expect the presence of holes in the initial configuration to significantly delay compression, even though this may increase the mixing time.

## 5.5 Achieving Compression and Expansion

In practice, Markov chain  $\mathcal{M}$  yields good compression. We simulated  $\mathcal{M}$  for  $\lambda = 4$  on 100 particles that began in a line; the configurations after 1, 2, 3, 4, and 5 million iterations of  $\mathcal{M}$  are shown in Figure 5.7. In contrast,  $\lambda = 2$ , while still favoring having more particle neighbors, does not yield compression; see Figure 5.8, where even after 20 million iterations of  $\mathcal{M}$ , the particles have not compressed. We conjecture there is a phase transition in  $\lambda$ , i.e., a critical value  $\lambda_c$  such that for all  $\lambda > \lambda_c$  the particles compress and for all  $\lambda < \lambda_c$  they do not compress. Such phase transitions exist for similar statistical physics models (e.g., [14]).

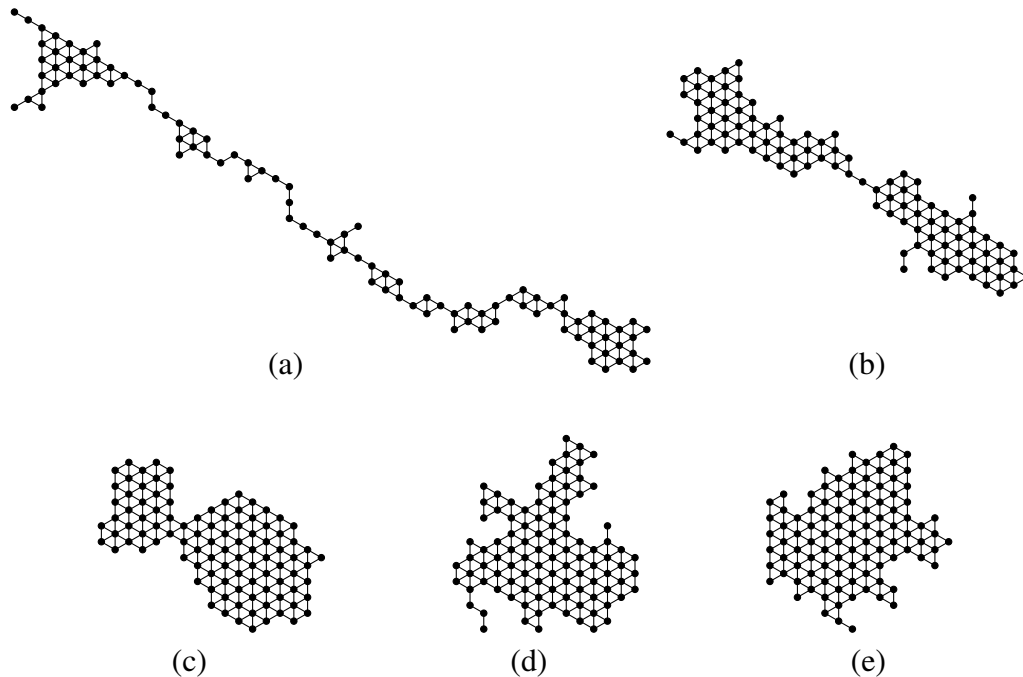


Figure 5.7: 100 particles in a line after (a) 1 million, (b) 2 million, (c) 3 million, (d) 4 million, and (e) 5 million iterations of  $\mathcal{M}$  with bias  $\lambda = 4$ . Edges have been drawn to indicate adjacencies between particles.

In this section we prove that these compression and expansion behaviors occur, though not all the way up to a critical point. Specifically, we show that for all  $\lambda > 2 + \sqrt{2}$ , there is a constant  $\alpha > 1$  such that at stationarity with all but exponentially small probability the particles are  $\alpha$ -compressed, meaning the perimeter of the system configuration is at




Figure 5.8: 100 particles in a line after (a) 10 million and (b) 20 million iterations of  $\mathcal{M}$  with bias  $\lambda = 2$ . Edges have been drawn to indicate adjacencies between particles.

most  $\alpha \cdot p_{min}$ , where  $p_{min}$  is the minimum possible perimeter of the particle system. We additionally prove that the same algorithm can be used for *expansion* for small values of  $\lambda$ ; in particular, for all  $0 < \lambda < 2.17$ , there is a constant  $\beta < 1$  such that at stationarity, with all but an exponentially small probability, the perimeter will be at least  $\beta \cdot p_{max}$ , where  $p_{max}$  is the maximum possible perimeter. This is counterintuitive because  $\lambda > 1$  corresponds to particles favoring having more neighbors, but this is not enough to guarantee compression.

### 5.5.1 Preliminaries: Counting Particle Configurations by Perimeter

Before we prove our main results, we first present some crucial lemmas focusing on counting the number of particle configurations with a given perimeter. Having good bounds on these numbers will strengthen our results about compression and expansion.

Let  $S_\alpha$  be the set of all hole-free particle configurations with perimeter at least  $\alpha \cdot p_{min}$  for some constant  $\alpha > 1$ , where  $p_{min}$  is the minimum possible perimeter for a configuration of  $n$  particles. We only consider hole-free configurations because we are concerned with behavior at stationarity and the stationary distribution  $\pi$  of  $\mathcal{M}$  only gives positive probability to hole-free configurations in  $\Omega^*$  (Corollary 5.19). We want an upper bound on  $\pi(S_\alpha) = \sum_{\sigma \in S_\alpha} \pi(\sigma)$ , the probability of being in a configuration with large perimeter, in

order to argue that this stationary probability of a configuration having large perimeter is exponentially small. Let  $c_k$  denote the number of hole-free configurations with perimeter  $k$  and recall that  $p_{max} = 2n - 2$  is the maximum possible perimeter for a configuration of  $n$  particles; using the expression for  $\pi$  given in Corollary 5.19, we can write  $\pi(S_\alpha)$  as:

$$\pi(S_\alpha) = \sum_{\sigma \in S_\alpha} \pi(\sigma) = \sum_{\sigma \in S_\alpha} \frac{\lambda^{-p(\sigma)}}{Z} = \frac{\sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{p_{max}} c_k \lambda^{-k}}{Z}.$$

Recall  $Z$  was defined in Corollary 5.19 as  $Z = \sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)}$ . In order to give an upper bound on  $\pi(S_\alpha)$ , we establish an upper bound on  $c_k$  and a lower bound on  $Z$ . The former we do in Lemma 5.23, and for the latter a trivial bound suffices for compression:  $Z \geq \lambda^{-p_{min}}$ .

To give an upper bound on the number  $c_k$  of hole-free configurations with perimeter  $k$ , we turn to lattice duality and self-avoiding walks; for a more thorough treatment of self-avoiding walks, see, e.g., [6].

**Definition 5.21.** *A self-avoiding walk (SAW) in a graph is a walk that never visits the same vertex twice.*

Self-avoiding walks are most commonly studied for graphs that are planar lattices, and we will focus on SAWs in the hexagonal lattice, also called the honeycomb lattice (Figure 5.9a). Examples of self-avoiding walks and non-self-avoiding walks in this lattice are shown in Figures 5.9b and 5.9c, respectively. The hexagonal lattice is of interest because it is dual to the triangular lattice  $\Gamma$  that particles occupy in our model. That is, by creating a new vertex in every face of the triangular lattice and connecting two of these new vertices if their corresponding triangular faces have a common edge, we obtain the hexagonal lattice; see Figure 5.10a.

The number of self-avoiding walks of a certain length from a given fixed starting point has been extensively studied for many planar lattices. This number is believed to grow exponentially with the length of the walk, and the base of this exponent is known as the *connective constant* of the lattice. More concretely, if  $N_l$  is the number of self-avoiding




Figure 5.9: (a) The hexagonal lattice. (b) A self-avoiding walk in the hexagonal lattice. (c) A walk that is not self-avoiding.




Figure 5.10: (a) The duality between the triangular lattice and the hexagonal lattice. (b) An example of a particle configuration  $\sigma$ , its corresponding polygon in the hexagonal lattice (shaded), and the self-avoiding walk bounding that polygon (bold).

walks of length  $l$  in some planar lattice  $L$ , then the *connective constant* of that lattice is defined as  $\mu_L = \lim_{l \rightarrow \infty} (N_l)^{1/l}$ . For example, for the square lattice  $2.625622 \leq \mu_{sq} \leq 2.679193$ , but an exact value has not been rigorously proved [70, 96]. The only lattice for which the connective constant is exactly known is our lattice of interest, the hexagonal lattice.

**Theorem 5.22** ([51]). *The  $l^{\text{th}}$  root of the number of self-avoiding walks of length  $l$  starting from a fixed vertex in the hexagonal lattice converges to  $\mu_{hex}$  as  $l \rightarrow \infty$ , where  $\mu_{hex} = \sqrt{2 + \sqrt{2}}$  is the connective constant of the hexagonal lattice.*

This theorem implies that the number of self-avoiding walks of length  $l$  in the hexagonal lattice is  $f(l) \cdot \mu_{hex}^l$ , for some subexponential function  $f$ .

To bound the number of hole-free particle configurations with some fixed perimeter, we turn from self-avoiding walks to the closely related notion of *self-avoiding polygons*,



where a self-avoiding polygon is a self-avoiding walk that starts and ends at the same vertex (Figure 5.10b). The number of self-avoiding walks of length  $l$  is an upper bound on the number of self-avoiding polygons of perimeter  $l$ .

**Lemma 5.23.** *The number of connected hole-free particle configurations with  $n$  particles and perimeter  $k$  is at most  $f(k) \cdot (2 + \sqrt{2})^k$  for some subexponential function  $f$ .*

*Proof.* To prove this theorem, we will consider the dual to the triangular lattice  $\Gamma$ , the hexagonal lattice  $\Gamma'$  (Figure 5.10a). For any connected particle configuration  $\sigma$  with  $n$  particles and no holes, consider the union  $A_\sigma$  of all the faces of  $\Gamma'$  corresponding to vertices of  $\Gamma$  that are occupied in  $\sigma$ . Whenever two particles are adjacent in  $\Gamma$ , their corresponding faces in  $\Gamma'$  share an edge. This union  $A_\sigma$  is a simply connected polygon because  $\sigma$  is connected and has no holes. The perimeter of  $A_\sigma$  is a self-avoiding polygon in the hexagonal lattice.

We next prove that if  $\sigma$  has perimeter  $k$ , then the perimeter of  $A_\sigma$  has  $2k + 6$  edges (equivalently,  $2k + 6$  vertices). We first note that a particle  $P$  is on the perimeter of  $\sigma$  if and only if its corresponding hexagon  $H_P$  in  $\Gamma'$  shares an edge with the perimeter of  $A_\sigma$ . That is, if a particle  $P$  appears on the perimeter of  $\sigma$  once with exterior angle  $\theta_P$ , where  $\theta_P \in \{120^\circ, 180^\circ, 240^\circ, 300^\circ, 360^\circ\}$ , then  $H_P$  has  $(\theta_P/60^\circ) - 1$  of its edges contained in the perimeter of  $A_\sigma$ . More generally, if a particle  $P$  appears on the perimeter  $m_P \geq 1$  times, with total exterior angles summing to  $\theta_P$ , then  $H_P$  has  $(\theta_P/60^\circ) - m_P$  of its edges contained in the perimeter of  $A_\sigma$ . For particle configurations  $\sigma$  with perimeter  $k$ , we conclude the number of edges on the perimeter of  $A_\sigma$  is:

$$p(A_\sigma) = \sum_{P \in p(\sigma)} \left( \frac{\theta_P}{60^\circ} - m_P \right) = \frac{1}{60} \left( \sum_{P \in p(\sigma)} \theta_P \right) - k = \frac{1}{60} (180k + 360) - k = 2k + 6.$$

The number of self-avoiding polygons of perimeter  $2k + 6$  in  $\Gamma'$  is an upper bound on the number of particle configurations with perimeter  $k$ . This value is itself less than the number of self-avoiding walks of length  $2k + 5$  in  $\Gamma'$ . As the connective constant for the

hexagonal lattice is  $\mu_{hex} = \sqrt{2 + \sqrt{2}}$ , there is some subexponential function  $f_1$  such that the number of these self-avoiding walks of length  $2k + 5$  is at most

$$f_1(2k + 5) \cdot (\mu_{hex})^{2k+5} = f_1(2k + 5) \cdot (2 + \sqrt{2})^{k+5/2}.$$

We conclude there is a subexponential function  $f(k) = f_1(2k + 5) \cdot \mu_{hex}^5$  such that the number of particle configurations with perimeter  $k$  is at most  $f(k) \cdot (2 + \sqrt{2})^k$ , the desired result.  $\square$

To prove our expansion results, we will also need some additional bounds on the number of particles configurations with certain perimeters. We will let  $S^\beta$  be the set of all hole-free configurations with perimeter at most  $\beta \cdot p_{max}$  for some constant  $0 < \beta < 1$ . Analogous to the approach for compression, we want to show  $\pi(S^\beta) = \sum_{\sigma \in S^\beta} \lambda^{-p(\sigma)} / Z$ , the stationary probability of being in a configuration with small perimeter, is exponentially small. The critical component of this result is an improved lower bound on  $Z$ . We give our first non-trivial lower bound on  $Z$  in Lemma 5.24, and this result is valid for all  $\lambda > 0$ . In Lemma 5.27 we get an improved lower bound on  $Z$  that is valid for all  $\lambda \geq 1$ . As above, obtaining these bounds requires a good bound on the number of configurations with certain perimeters.

We do this by giving a lower bound on the number of configurations with  $n$  particles and a given perimeter (as opposed to the previous lemma, where we found upper bounds for this quantity). To begin, we recall  $p_{max} = 2n - 2$  and note:

$$Z = \sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)} \geq \sum_{\substack{\sigma \in \Omega^* : \\ p(\sigma) = p_{max}}} \lambda^{-p(\sigma)} = c_{2n-2} \lambda^{-(2n-2)},$$

where  $c_{2n-2}$  is the number of configurations with perimeter exactly  $2n - 2$ . Note if a configuration  $\sigma$  with  $n$  particles has perimeter  $2n - 2$ , then by Lemmas 5.3 and 5.4 it must

be that  $\sigma$  has exactly  $n - 1$  edges and no triangles; that is,  $\sigma$  is an induced tree in  $\Gamma$ . We present a method for enumerating a subset of these trees, giving a lower bound on  $c_{2n-2}$ .

**Lemma 5.24.** *For any  $\lambda > 0$ , it holds that  $Z \geq (\sqrt{2}/\lambda)^{p_{max}}$ .*

*Proof.* We enumerate  $n$ -vertex paths in  $\Gamma$  where every step is either down-right or up-right; this is a subset of the trees contributing to  $c_{2n-2}$ . Starting from the first particle, there are  $2^{n-1}$  ways to place rest of the particles to form such a path, where each one is either up-right or down-right from the previous one. This means there are at least  $2^{n-1}$  such paths, giving

$$c_{2n-2} \geq 2^{n-1} = \sqrt{2}^{2n-2} = \sqrt{2}^{p_{max}}.$$

From this, it follows that

$$Z = \sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)} \geq \sum_{\substack{\sigma \in \Omega^* : \\ p(\sigma) = p_{max}}} \lambda^{-p_{max}} \geq \sqrt{2}^{p_{max}} \lambda^{-p_{max}} = \left( \frac{\sqrt{2}}{\lambda} \right)^{p_{max}}.$$

□

This bound could be improved significantly with a better lower bound for  $c_{2n-2}$ , though this is eclipsed by the lower bound for  $Z$  when  $\lambda \geq 1$  given next. The key observation is that for  $\lambda > 1$ , for any value  $k < 2n - 2$ ,  $\lambda^{-k} > \lambda^{-(2n-2)}$ . Thus as  $p_{max} = 2n - 2$ , it follows that

$$Z = \sum_{\sigma} \lambda^{-p(\sigma)} \geq \sum_{\sigma} \lambda^{-(2n-2)}.$$

Thus it suffices to find a lower bound on the total number of connected, hole-free configurations with  $n$  particles and any perimeter, instead of only counting the number of configurations with maximum perimeter. Exploiting this, we will be able to get a better lower bound on  $Z$  than in the case above where  $\lambda$  was unrestricted.

**Lemma 5.25.** *If  $\lambda \geq 1$ , then  $Z \geq 0.12 \cdot (1.67/\lambda)^{p_{max}}$ .*




Figure 5.11: All 11 connected hole-free configurations with three particles. In each the highest leftmost particle is labeled  $H$ , and the lowest leftmost particle is labeled  $L$ ; when there is only one leftmost particle  $H = L$ .

*Proof.* We give a lower bound on the number of connected, hole-free configurations on  $n$  particles by iteratively enumerating a subset of them. Note there are 11 connected hole-free configurations with exactly 3 particles; all 11 are shown in Figure 5.11.

Given some hole-free configuration  $\sigma$  with  $1 + 3j$  particles,  $j \geq 0$ , we show how to enumerate  $22$  distinct hole-free configurations of  $4 + 3j$  particles. Let  $P$  be the highest rightmost particle of  $\sigma$  and let  $Q$  be the lowest rightmost particle of  $\sigma$ ; possibly  $P = Q$ . Choose any of the 11 hole-free configurations with 3 particles, and let  $L$  be its lowest leftmost particle and  $H$  be its highest leftmost particle as in Figure 5.11; possibly  $H = L$ . Attach this configuration to  $\sigma$  either by placing  $H$  below and right of  $Q$  or by placing  $L$  above and right of  $P$ ; see Figure 5.12 for two such examples. Note even if  $Q = P$  and  $H = L$ , this still results in two distinct configurations. In the first case, all locations directly below  $Q$  and all locations directly above  $H$  are unoccupied; this ensures the only adjacency between  $\sigma$  and the newly added three particles is between  $Q$  and  $H$ , meaning no holes have been created. Similarly in the second case, all locations above  $P$  or below  $L$  are unoccupied, again ensuring no holes form.

Using this process and beginning with a single particle (as in Figure 5.12a), we can enumerate  $22^j$  distinct configurations with  $1 + 3j$  particles for all  $j \geq 0$ . This does not enumerate all configurations on  $1 + 3j$  particles: for example, there are 42 configurations

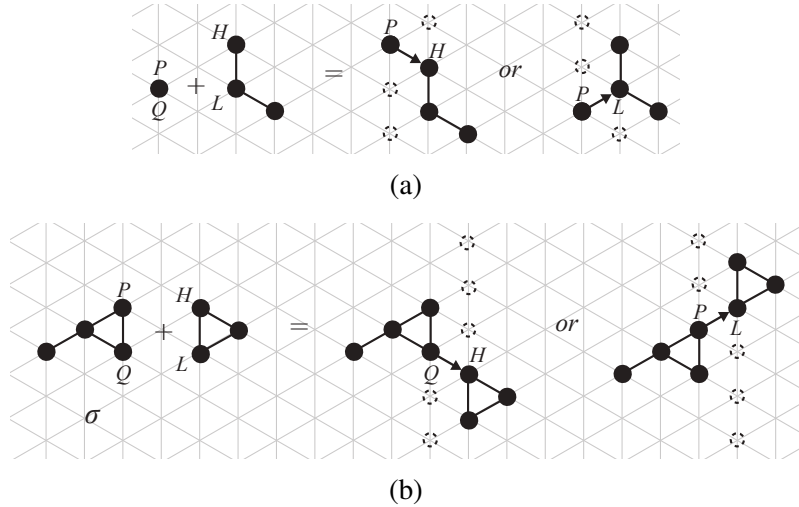


Figure 5.12: The iterative process of Lemma 5.25. (a) One of the 11 connected hole-free configurations with three particles, and the two ways it can attach to the single particle with which the iterative process begins. (b) Another of the 11 connected hole-free configurations on three vertices, and the two ways it can attach to a configuration  $\sigma$  with four particles. Particle adjacencies have been drawn as black lines, and dashed circles indicate unoccupied locations that guarantee no hole exists in the constructed configurations.

on 4 particles and this process only enumerates 22 of them. However, this process iterates nicely and produces reasonable lower bounds as the number of particles gets large.

To get a lower bound on the number of configurations of  $n$  particles when  $n \not\equiv 1 \pmod{3}$ , we can simply enumerate all configurations on  $1 + 3j \leq n$  particles for  $j = \lfloor \frac{n-1}{3} \rfloor$ , and add one or two particles to each in some deterministic way. We conclude that for any  $n$ , the number of hole-free connected configurations of  $n$  particles is at least

$$22^{\lfloor \frac{n-1}{3} \rfloor} \geq 22^{\frac{n-1}{3}} \cdot 22^{-2/3} = 22^{-2/3} (22^{1/6})^{2n-2} > 0.12 \cdot 1.67^{2n-2}.$$

Using this bound, it follows that

$$Z = \sum_{\sigma} \lambda^{-p(\sigma)} \geq \sum_{\sigma} \lambda^{-(2n-2)} > 0.12 \cdot 1.67^{2n-2} \cdot \lambda^{-(2n-2)} = 0.12 \cdot \left( \frac{1.67}{\lambda} \right)^{p_{max}}.$$

□

This bound can be improved even further by iteratively adding hole-free configurations

of 50 particles instead of three particles. A result of Jensen [69] will be essential. In that paper, the author presents a parallel algorithm efficient enough to count the number of benzenoid hydrocarbons containing  $h$  hexagonal cells up to  $h = 50$ . A benzenoid hydrocarbon containing  $h$  hexagonal cells is exactly equivalent to a connected particle configuration with no holes and  $h$  particles, implying the next lemma.

**Lemma 5.26** ([69]). *The number of connected particle configurations with no holes and 50 particles is*

$$N_{50} = 2,430,068,453,031,180,290,203,185,942,420,933.$$

**Lemma 5.27.** *If  $\lambda \geq 1$ , then  $Z \geq 0.13 \cdot (2.17/\lambda)^{p_{max}}$ .*

*Proof.* We use the same approach as in Lemma 5.25, noting that  $2.17 \sim (2N_{50})^{1/100}$ . To get a lower bound on the number of configurations with  $n$  particles, first write  $n$  as  $n = 1 + 50i + j$ , where  $i, j \in \mathbb{Z}_{\geq 0}$  and  $j < 50$ ; subject to these requirements,  $i$  and  $j$  are unique. Iteratively construct one particle configuration  $\sigma$  with  $n$  particles by beginning with a single particle and repeatedly attaching one of the  $N_{50}$  configurations with 50 particles to the right as in the proof of Lemma 5.25: place its highest leftmost particle  $H$  below and right of the existing configuration's lowest rightmost particle  $Q$ , or place its lowest leftmost particle  $L$  above and right of the existing configuration's highest rightmost particle  $P$ . This process, applied  $i$  times, yields a hole-free configuration on  $1 + 50i = n - j$  particles. There are then  $2N_j$  ways, following the same procedure, to attach the remaining  $j$  particles to form a hole-free configuration with  $n$  particles. In this way, we can enumerate  $(2N_{50})^i \cdot 2N_j$  unique hole-free configurations on  $n$  particles. It follows that the number of connected hole-free configurations on  $n$  particles is at least

$$(2N_{50})^i \cdot 2N_j = (2N_{50})^{\frac{n-1-j}{50}} \cdot 2N_j = (2N_{50})^{\frac{n-1}{50}} \cdot (2N_{50})^{-\frac{j}{50}} \cdot 2N_j.$$

Calculations show that for all  $0 \leq j < 50$ ,  $(2N_{50})^{-j/50} \cdot 2N_j \geq 0.13$ . It follows that the

number of connected hole-free configurations with  $n$  particles is at least

$$0.13 \cdot \left( (2N_{50})^{1/100} \right)^{2n-2} = 0.13 \cdot \left( (2N_{50})^{1/100} \right)^{p_{max}}.$$

Noting that  $(2N_{50})^{1/100} > 2.17$ , it follows that

$$Z \geq \sum_{\sigma} \lambda^{-(2n-2)} \geq 0.13 \cdot (2.17)^{p_{max}} \lambda^{-p_{max}} = 0.13 \cdot (2.17/\lambda)^{p_{max}}.$$

□

As we will see in Section 5.5.3, this will directly imply that the particle system will not exhibit compression for any  $\lambda < 2.17$ . We expect this bound will improve given accurate counts of the number of particle configurations for even larger  $n$ . Computationally this seems infeasible, and a careful analysis of the work done in [69] suggests the best bound achievable by this method would be expansion for all  $\lambda < 2.27$ , only a mild improvement and still far from the known lower bound for compression,  $\lambda > 2 + \sqrt{2}$ .

### 5.5.2 Achieving Compression

We proved in Section 5.4.6 that Markov chain  $\mathcal{M}$  converges to a unique stationary distribution, and we know that distribution exactly (Corollary 5.19). In this section, we show that when parameter  $\lambda$  is large enough, this stationary distribution exhibits compression with high probability. While we conjecture compression occurs even earlier, before  $\mathcal{M}$  is close to stationarity, our proofs rely on analyzing the stationary distribution of  $\mathcal{M}$ .

Recall for any  $\alpha > 1$  we say a configuration  $\sigma$  with  $n$  particles is  $\alpha$ -compressed if its perimeter  $p(\sigma) < \alpha \cdot p_{min}$ , where  $p_{min}$  is the minimum possible perimeter of a configuration with  $n$  particles. We prove that, for any  $\alpha > 1$  and provided  $\lambda$  and  $n$  are large enough, a configuration chosen at random according to the stationary distribution of  $\mathcal{M}$  is  $\alpha$ -compressed with all but a probability that is exponentially small (in  $n$ ). Values of  $\alpha$  closer to 1 simply require larger  $\lambda$  values. Conversely, we then prove (as a corollary) that

for any  $\lambda > 2 + \sqrt{2}$ , there is a constant  $\alpha$  such that with high probability  $\alpha$ -compression occurs at stationarity.

To simplify notation, we define the *weight* of a configuration  $\sigma$  to be  $w(\sigma) = \pi(\sigma) \cdot Z = \lambda^{-p(\sigma)}$ . For a set  $S \subseteq \Omega$ , we define  $w(S)$  as the sum of the weights of all configurations in  $S$ . We now prove our main result.

**Theorem 5.28.** *For any  $\alpha > 1$ , let  $\lambda^* = (2 + \sqrt{2})^{\frac{\alpha}{\alpha-1}}$ . There exists  $n^* \geq 0$  and  $\gamma < 1$  such that for all  $\lambda > \lambda^*$  and  $n > n^*$ , the probability that a random sample  $\sigma$  drawn according to the stationary distribution  $\pi$  of  $\mathcal{M}$  is not  $\alpha$ -compressed is exponentially small:*

$$\mathbb{P}_{\sigma \sim \pi} (p(\sigma) \geq \alpha \cdot p_{min}) < \gamma^{\sqrt{n}}.$$

*Proof.* Let  $S_\alpha$  be the set of configurations of perimeter at least  $\alpha \cdot p_{min}$ . We wish to show that  $\pi(S_\alpha)$  is smaller than some function that is exponentially small in  $n$ .

We first consider the normalizing constant  $Z$  of  $\pi$ ; recall  $Z = \sum_{\sigma \in \Omega^*} \lambda^{-p(\sigma)}$ . If  $\sigma_{min}$  is a configuration of  $n$  particles achieving the minimum possible perimeter  $p_{min}$ , then  $w(\sigma_{min}) = \lambda^{-p_{min}}$  is a lower bound on  $Z$ . It follows that

$$\pi(S_\alpha) = \frac{w(S_\alpha)}{Z} < \frac{w(S_\alpha)}{w(\sigma_{min})}.$$

The remainder of this proof will be spent finding an upper bound on  $w(S_\alpha)/w(\sigma_{min})$  that is exponentially small in  $n$ . To begin, we stratify  $S_\alpha$  into sets of configurations that have the same perimeter. Let  $A_k$  be the set of all configurations with perimeter  $k \in \mathbb{Z}$ ; then  $S_\alpha = \bigcup_{k=\lceil \alpha \cdot p_{min} \rceil}^{p_{max}} A_k$ . We can then write

$$\frac{w(S_\alpha)}{w(\sigma_{min})} \leq \frac{\sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} w(A_k)}{\lambda^{-p_{min}}}.$$

The weight of each element in the set  $A_k$  is the same,  $\lambda^{-k}$ . Letting  $\nu := 2 + \sqrt{2} \sim 3.42$  for simplicity, by Lemma 5.23 the number of elements in set  $A_k$  is at most  $f(k)\nu^k$  for some



subexponential function  $f$ . We see that

$$\frac{w(S_\alpha)}{w(\sigma_{min})} \leq \frac{\sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} f(k) \nu^k \lambda^{-k}}{\lambda^{-p_{min}}} = \sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} f(k) \nu^{(1-\log_\nu \lambda)k + (\log_\nu \lambda)p_{min}}.$$

Using the inequality  $p_{min} \leq k/\alpha$  and noting that  $\lambda > \nu$ , it follows that

$$\frac{w(S_\alpha)}{w(\sigma_{min})} \leq \sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} f(k) \nu^{(1-\log_\nu \lambda)k + (\log_\nu \lambda)(k/\alpha)} = \sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} f(k) \nu^{(1-(1-1/\alpha)\log_\nu \lambda)k}.$$

As  $\lambda > \lambda^* = \nu^{\frac{\alpha}{\alpha-1}}$ , it follows that  $\log_\nu \lambda > \frac{\alpha}{\alpha-1}$ . We then see that the exponent in the previous expression satisfies

$$1 - \left(1 - \frac{1}{\alpha}\right) \log_\nu \lambda < 1 - \frac{\alpha-1}{\alpha} \cdot \frac{\alpha}{\alpha-1} = 0.$$

For simplicity, we let  $c$  be the negation of this expression,  $c := -1 + (1 - 1/\alpha) \log_\nu \lambda > 0$ .

By applying the inequality  $k \geq \alpha \cdot p_{min}$  and using Lemma 5.1 which states that  $p_{min} > \sqrt{n}$ , we see that

$$\frac{w(S_\alpha)}{w(\sigma_{min})} \leq \sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} f(k) \nu^{-ck} \leq \sum_{k=\lceil \alpha \cdot p_{min} \rceil}^{2n-2} f(k) \nu^{-c\alpha \sqrt{n}}.$$

Let  $f_1(n) = \sum_{k=\sqrt{n}}^{2n-2} f(k)$ ; we note that this is also a function that is subexponential in  $n$  as it is a sum of a linear number of subexponential terms. We can then say that

$$\frac{w(S_\alpha)}{w(\sigma_{min})} \leq f_1(n) (\nu^{-c\alpha})^{\sqrt{n}}.$$

As  $c > 0$  and  $f_1$  is subexponential, we conclude that because  $\nu^{-c\alpha} < 1$  there exists  $\gamma$  satisfying  $\nu^{-c\alpha} < \gamma < 1$  and an  $n^*$  such that for all  $n \geq n^*$ ,

$$\mathbb{P}_{\sigma \sim \pi}(p(\sigma) \geq \alpha \cdot p_{min}) = \pi(S_\alpha) \leq \frac{w(S_\alpha)}{w(\sigma_{min})} \leq f_1(n) (\nu^{-c\alpha})^{\sqrt{n}} < \gamma^{\sqrt{n}}.$$

□

While the above result shows that  $\mathcal{M}$  accomplishes  $\alpha$ -compression for any  $\alpha > 1$ , the smaller we want  $\alpha$  to be the larger  $\lambda$  needs to be. In practice, when  $\lambda$  is large  $\mathcal{M}$  takes a very long time to reach any compressed configuration. Because of this, what happens when  $\lambda$  is small is also of interest. We now show that provided  $\lambda > 2 + \sqrt{2}$  there is some constant  $\alpha$  such that  $\alpha$ -compression occurs. Of course, there is again a trade-off: the smaller  $\lambda$  is, the larger  $\alpha$  is.

**Corollary 5.29.** *For any  $\lambda > 2 + \sqrt{2} =: \nu$ , for any constant  $\alpha > \log_\nu \lambda / (\log_\nu \lambda - 1)$  there exists  $n^* \geq 0$  and  $\gamma < 1$  such that for all  $n \geq n^*$ , a random sample  $\sigma$  drawn according to the stationary distribution  $\pi$  of  $\mathcal{M}$  satisfies*

$$\mathbb{P}_{\sigma \sim \pi} (p(\sigma) \geq \alpha \cdot p_{min}) < \gamma^{\sqrt{n}}.$$

*Proof.* If  $\alpha > \frac{\log_\nu \lambda}{\log_\nu \lambda - 1}$ , then solving for  $\lambda$  gives  $\lambda > \nu^{\frac{\alpha}{\alpha-1}}$ . Theorem 5.28 then gives the desired result.  $\square$

### 5.5.3 Using $\mathcal{M}$ for Expansion

A nice feature of our algorithm is that it also provably achieves particle expansion for different values of bias parameter  $\lambda$ . Analogous to our definition of  $\alpha$ -compression, we say a configuration  $\sigma$  is  $\beta$ -expanded for some  $0 < \beta < 1$  if  $p(\sigma) > \beta \cdot p_{max}$ , where  $p_{max} = 2n - 2$ .

For a configuration of  $n$  particles,  $p_{max} = \Theta(n)$  and  $p_{min} = \Theta(\sqrt{n})$ , so  $\beta$ -expansion and  $\alpha$ -compression for any constants  $\beta$  and  $\alpha$  are mutually exclusive for sufficiently large  $n$ . We prove in this section that, for all  $0 < \lambda < 2.17$  and provided  $n$  is large enough, there is a constant  $\beta$  such that a configuration chosen at random according to the stationary distribution of  $\mathcal{M}$  is  $\beta$ -expanded with all but exponentially small probability. This is notable because it implies that, counter-intuitively,  $\lambda > 1$  (i.e., favoring more neighbors) is not sufficient to guarantee particle compression as one might first guess.

We first show that for any value of  $\beta$  it is possible to achieve  $\beta$ -expansion by simply running  $\mathcal{M}$  with input parameter  $\lambda$  sufficiently small. The closer  $\beta$  is to 1, the closer  $\lambda$  must be to 0 in order to achieve  $\beta$ -expansion.

**Theorem 5.30.** *For any  $0 < \beta < 1$ , let  $\lambda^* = (\sqrt{2}/(2 + \sqrt{2})^\beta)^{\frac{1}{1-\beta}}$ . There exists  $n^* \geq 0$  and  $\gamma < 1$  such that for all  $\lambda < \lambda^*$  and  $n \geq n^*$ , the probability that a random sample  $\sigma$  drawn according to the stationary distribution  $\pi$  of  $\mathcal{M}$  is not  $\beta$ -expanded is exponentially small:*

$$\mathbb{P}_{\sigma \sim \pi}(p(\sigma) \leq \beta \cdot p_{max}) < \gamma^{\sqrt{n}}$$

*Proof.* We begin by rewriting  $\lambda^*$  in a way that will be convenient:

$$\lambda^* = \left( \frac{\sqrt{2}}{(2 + \sqrt{2})^\beta} \right)^{\frac{1}{1-\beta}} = \left( \sqrt{2} \right)^{\frac{1}{1-\beta} \log_{\sqrt{2}} \left( \frac{\sqrt{2}}{(2 + \sqrt{2})^\beta} \right)} = \left( \sqrt{2} \right)^{\frac{1 - \beta \log_{\sqrt{2}}(2 + \sqrt{2})}{1-\beta}}.$$

As  $\log_{\sqrt{2}}(2 + \sqrt{2}) \sim 1.356 > 1$  and  $0 < \beta < 1$ , then  $1 - \beta \log_{\sqrt{2}}(2 + \sqrt{2}) < 1 - \beta$  so the exponent in the above expression is always less than one. This means that  $\lambda^* < \sqrt{2}$ , and thus  $\lambda < \sqrt{2}$ , a fact we will use later.

Let  $S^\beta$  be the set of configurations of perimeter at most  $\beta \cdot p_{max}$ . We wish to show that  $\pi(S^\beta)$  is smaller than some function that is exponentially small in  $n$ .

Applying Lemma 5.24, which gives an upper bound on the normalizing constant  $Z$  of stationary distribution  $\pi$ , we see that

$$\pi(S^\beta) = \frac{w(S^\beta)}{Z} \leq \frac{w(S^\beta)}{\left(\frac{\sqrt{2}}{\lambda}\right)^{p_{max}}} = w(S^\beta) \left(\frac{\lambda}{\sqrt{2}}\right)^{p_{max}}.$$

The remainder of this proof will be spent finding an upper bound on the right hand side of the above equation that is exponentially small in  $n$ . To begin, we stratify  $S^\beta$  into sets of configurations that have the same perimeter. Let  $B_k$  be the set of all configurations with

perimeter  $k$ ; then  $S^\beta = \bigcup_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} B_k$ . We can write

$$w(S^\beta) \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}} = \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} w(B_k) \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}}.$$

The weight of each element in the set  $B_k$  is the same,  $\lambda^{-k}$ . By Lemma 5.23, the number of elements in set  $B_k$  is at most  $f(k)(2 + \sqrt{2})^k$  for some subexponential function  $f$ . We see that

$$w(S^\beta) \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}} \leq \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} f(k)(2 + \sqrt{2})^k \lambda^{-k} \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}}.$$

Recalling that  $k \leq \beta \cdot p_{\max}$ , meaning  $p_{\max} \geq k/\beta$ , then as  $\lambda < \sqrt{2}$  we see that

$$\begin{aligned} w(S^\beta) \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}} &\leq \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} f(k)(2 + \sqrt{2})^k \lambda^{-k} \left( \frac{\lambda}{\sqrt{2}} \right)^{\frac{k}{\beta}} \\ &= \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} f(k) \left( \frac{2 + \sqrt{2}}{\lambda} \cdot \left( \frac{\lambda}{\sqrt{2}} \right)^{\frac{1}{\beta}} \right)^k. \end{aligned}$$

Let  $\xi = \frac{2 + \sqrt{2}}{\lambda} \cdot \left( \frac{\lambda}{\sqrt{2}} \right)^{\frac{1}{\beta}}$ ; we see that as  $\lambda < \lambda^*$ , then

$$\xi = \frac{2 + \sqrt{2}}{\lambda} \cdot \left( \frac{\lambda}{\sqrt{2}} \right)^{\frac{1}{\beta}} = \frac{2 + \sqrt{2}}{\sqrt{2}^{1/\beta}} \cdot \lambda^{\frac{1-\beta}{\beta}} < \frac{2 + \sqrt{2}}{\sqrt{2}^{1/\beta}} \cdot \left( \sqrt{2}/(2 + \sqrt{2}) \right)^{\frac{1}{\beta}} = 1.$$

By applying the inequality  $k \geq p_{\min}$  and recalling that  $p_{\min} > \sqrt{n}$  (Lemma 5.1), we see that

$$w(S^\beta) \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}} \leq \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} f(k)\xi^k < \sum_{k=\sqrt{n}}^{2n-2} f(k)\xi^{\sqrt{n}}.$$

Let  $f_1(n) = \sum_{k=\sqrt{n}}^{2n-2} f(k)$ . Because  $f_1$  has at most  $2n$  summands and  $f$  is a subexponential function, then  $f_1(n)$  is a subexponential function. We conclude there exists a  $\gamma$  such that  $\xi < \gamma < 1$  and an  $n^*$  such that for all  $n \geq n^*$ ,

$$\mathbb{P}_{\sigma \sim \pi}(p(\sigma) \leq \beta \cdot p_{\max}) = \pi(S^\beta) \leq w(S^\beta) \left( \frac{\lambda}{\sqrt{2}} \right)^{p_{\max}} < f_1(n)\xi^{\sqrt{n}} < \gamma^{\sqrt{n}}.$$

□

While the above result shows that  $\mathcal{M}$  accomplishes  $\beta$ -expansion for any  $\beta < 1$ , the larger we want  $\beta$  to be the smaller  $\lambda$  needs to be. However, larger values of  $\lambda$  are still of interest. Just as for compression, we now show that provided  $\lambda < \sqrt{2}$  there is some constant  $\beta$  such that  $\beta$ -expansion occurs. Of course, there is again a trade-off: the larger  $\lambda$  is, the smaller  $\beta$  is.

**Corollary 5.31.** *For all  $0 < \lambda < \sqrt{2}$ , for any constant  $\beta < \frac{\log_{\sqrt{2}} \lambda - 1}{\log_{\sqrt{2}} \lambda - \log_{\sqrt{2}}(2 + \sqrt{2})}$ , there exists  $n^* \geq 0$  and  $\gamma < 1$  such that for all  $n \geq n^*$ , a random sample  $\sigma$  drawn according to the stationary distribution  $\pi$  of  $\mathcal{M}$  satisfies*

$$\mathbb{P}_{\sigma \sim \pi} (p(\sigma) \leq \beta \cdot p_{max}) < \gamma^{\sqrt{n}}.$$

*Proof.* Theorem 5.30 holds whenever  $\lambda < \left( \frac{\sqrt{2}}{(2 + \sqrt{2})^\beta} \right)^{\frac{1}{1-\beta}}$ . Solving for  $\beta$ , we see the theorem applies whenever  $\beta < \frac{\log_{\sqrt{2}} \lambda - 1}{\log_{\sqrt{2}} \lambda - \log_{\sqrt{2}}(2 + \sqrt{2})}$ , as desired. □

When we know  $\lambda \geq 1$ , the improved bounds in Lemma 5.27 can be used to show  $\beta$ -expansion occurs for an even greater range of values for  $\lambda$ . Again, larger values of  $\lambda$  necessitate smaller, but still constant, values of  $\beta$ .

**Theorem 5.32.** *For all  $1 \leq \lambda < x := (2N_{50})^{1/100} \sim 2.17$ , for any  $\beta < \frac{\log_x \lambda - 1}{\log_x \lambda - \log_x(2 + \sqrt{2})}$  there exists  $n^* \geq 0$  and  $\gamma < 1$  such that for all  $n \geq n^*$ , a random sample  $\sigma$  drawn according to the stationary distribution  $\pi$  of  $\mathcal{M}$  satisfies*

$$\mathbb{P}_{\sigma \sim \pi} (p(\sigma) \leq \beta \cdot p_{max}) < \gamma^{\sqrt{n}}.$$

*Proof.* Let  $S^\beta$  be the set of configurations of perimeter at most  $\beta \cdot p_{max}$ . We wish to show that  $\pi(S^\beta)$  is smaller than some function that is exponentially small in  $n$ .

Applying Lemma 5.27, which gives an upper bound on the normalizing constant  $Z$  of

stationary distribution  $\pi$ , we see that

$$\pi(S^\beta) = \frac{w(S^\beta)}{Z} \leq \frac{w(S^\beta)}{0.13 \left(\frac{x}{\lambda}\right)^{p_{max}}} \leq 8 w(S^\beta) \left(\frac{\lambda}{x}\right)^{p_{max}}.$$

The remainder of this proof will be spent finding an upper bound on the right hand side of the above equation that is exponentially small in  $n$ . To begin, we stratify  $S^\beta$  into sets of configurations that have the same perimeter. Let  $B_k$  be the set of all configurations with perimeter  $k$ ; then  $S^\beta = \bigcup_{k=p_{min}}^{\lfloor \beta \cdot p_{max} \rfloor} B_k$ . We can then write

$$8 w(S^\beta) \left(\frac{\lambda}{x}\right)^{p_{max}} = \sum_{k=p_{min}}^{\lfloor \beta \cdot p_{max} \rfloor} 8 w(B_k) \left(\frac{\lambda}{x}\right)^{p_{max}}$$

The weight of each element in the set  $B_k$  is the same,  $\lambda^{-k}$ . By Lemma 5.23, the number of elements in set  $B_k$  is at most  $f(k)(2 + \sqrt{2})^k$  for some subexponential function  $f$ . We see that

$$8 w(S^\beta) \left(\frac{\lambda}{x}\right)^{p_{max}} \leq \sum_{k=p_{min}}^{\lfloor \beta \cdot p_{max} \rfloor} 8 f(k)(2 + \sqrt{2})^k \lambda^{-k} \left(\frac{\lambda}{x}\right)^{p_{max}}.$$

Recalling that  $k < \beta \cdot p_{max}$ , meaning  $p_{max} \geq k/\beta$ , then as  $\lambda < x$  we see that

$$\begin{aligned} 8 w(S^\beta) \left(\frac{\lambda}{x}\right)^{p_{max}} &\leq \sum_{k=p_{min}}^{\lfloor \beta \cdot p_{max} \rfloor} 8 f(k)(2 + \sqrt{2})^k \lambda^{-k} \left(\frac{\lambda}{x}\right)^{\frac{k}{\beta}} \\ &= \sum_{k=p_{min}}^{\lfloor \beta \cdot p_{max} \rfloor} 8 f(k) \left(\frac{2 + \sqrt{2}}{\lambda} \cdot \left(\frac{\lambda}{x}\right)^{\frac{1}{\beta}}\right)^k. \end{aligned}$$

Let  $\xi = \frac{2 + \sqrt{2}}{\lambda} \cdot \left(\frac{\lambda}{x}\right)^{\frac{1}{\beta}}$ . The condition  $\beta < \frac{\log_x \lambda - 1}{\log_x \lambda - \log_x (2 + \sqrt{2})}$  can be equivalently expressed as  $\lambda < (x/(2 + \sqrt{2})^\beta)^{\frac{1}{1-\beta}}$ . It follows that

$$\xi = \frac{2 + \sqrt{2}}{\lambda} \cdot \left(\frac{\lambda}{x}\right)^{\frac{1}{\beta}} = \frac{2 + \sqrt{2}}{x^{1/\beta}} \cdot \lambda^{\frac{1-\beta}{\beta}} < \frac{2 + \sqrt{2}}{x^{1/\beta}} \cdot \left(\frac{x}{(2 + \sqrt{2})^\beta}\right)^{\frac{1}{\beta}} = 1.$$

By applying the inequality  $k \geq p_{min}$  and recalling that  $p_{min} > \sqrt{n}$  (Lemma 5.1), we see

that

$$8 w(S^\beta) \left(\frac{\lambda}{x}\right)^{p_{max}} \leq \sum_{k=p_{min}}^{\lfloor \beta \cdot p_{max} \rfloor} 8 f(k) \xi^k < \sum_{k=\sqrt{n}}^{2n-2} 8 f(k) \xi^{\sqrt{n}}.$$

Let  $f_1(n) = \sum_{k=\sqrt{n}}^{2n-2} 8 f(k)$ . Because  $f_1$  has at most  $2n$  summands and  $f$  is a subexponential function, then  $f_1(n)$  is a subexponential function. We conclude there exists a  $\gamma$  such that  $\xi < \gamma < 1$  and an  $n^*$  such that for all  $n \geq n^*$ ,

$$\mathbb{P}_{\sigma \sim \pi}(p(\sigma) \leq \beta \cdot p_{max}) = \pi(S^\beta) \leq w(S^\beta) \left(\frac{\lambda}{x}\right)^{p_{max}} < f_1(n) \xi^{\sqrt{n}} < \gamma^{\sqrt{n}}$$

□

Combining Theorem 5.32 with Corollary 5.31 gives the following result.

**Corollary 5.33.** *For all  $0 < \lambda < 2.17$ , there exists a constant  $0 < \beta < 1$  such that with all but exponentially small probability a sample drawn according to stationary distribution  $\pi$  of  $\mathcal{M}$  is  $\beta$ -expanded.*

*Proof.* By Corollary 5.31, this is true for  $0 < \lambda < \sqrt{2}$ . By Theorem 5.32, this is true for  $1 \leq \lambda < 2.17$ . □

## 5.6 Shortcut Bridging: Overview, Problem Definition, and Preliminaries

In this and subsequent sections, we consider a very different application of the stochastic approach to algorithms for programmable matter: *shortcut bridging*. This work demonstrates that many fundamental elements of this stochastic approach can be generalized to applications beyond the specific context of compression.

We present an algorithm inspired by the work of Reid et al. [98], who found that army ants continuously modify the shape and position of foraging bridges — constructed and maintained by their own bodies — across holes and uneven surfaces in the forest floor. These bridges appear to stabilize in a structural formation that balances the “benefit of increased foraging trail efficiency” with the “cost of removing workers from the foraging pool to form the structure” [98]. We attempt to capture this inherent trade-off in our

algorithm for “shortcut bridging” in self-organizing particle systems (formally defined in Section 5.6.2). Shortcut bridging is an attractive goal for programmable matter systems, as many application domains envision deploying programmable matter on surfaces with structural irregularities or dynamic topologies. For example, one commonly imagined application of smart sensor networks is to detect and span small cracks in infrastructure such as roads or bridges; dynamic bridging behavior would enable the system to remain connected and shift position as cracks form.

Using the same Markov chain techniques as above, we rigorously analyze our algorithm, show it achieves a near-optimal balance between the competing factors of path length and bridge cost, and prove that it exhibits a dependence on the angle of the gap being “shortcut” similar to that of the ant bridges. We also present simulation results that qualitatively compare our algorithm with the army ant bridging behavior. Our work gives a plausible explanation of how convergence to globally optimal configurations can be achieved via local interactions by simple organisms (e.g., ants) with some limited computational power and access to random bits. The proposed algorithm also demonstrates the robustness of the stochastic approach to algorithms for programmable matter.

### 5.6.1 Related Work

The work of Reid et al. [98] showing that army ants of the genus *Eciton* can build bridges whose structure seems to optimize a global trade-off is just one example of a simple natural system without centralized control finding a solution to an optimization problem. For example, single-celled slime molds have famously exhibited an ability to find solutions to mazes [95] and, more recently, to solve multi-arm bandit problems [99]. The phenomenon of local interactions yielding emergent, collective behavior in natural systems comprised of many individuals is also more general than just the ant behavior observed in [98]. For example, honey bees choose hive locations based on decentralized recruitment [19]; fire ants, when their nests are flooded, form floating rafts out of their own bodies [89]; and cockroach



larvae perform self-organizing aggregation using pheromones with limited range [68].

In the molecular programming domain, simpler variations of bridging have been studied. Groundbreaking works in this area, such as that of Mohammed et al. [90], focus on forming molecular structures that connect some fixed points; our work may offer insights on further optimizing the quality and/or cost of the resulting bridges.

### 5.6.2 Problem Definition

Just as the uneven surfaces of the forest floor affect the foraging behavior of army ants, the collective behavior of particle systems should change when  $\Gamma$  is non-uniform. Here, we focus on system behaviors when the vertices of  $\Gamma$  are either *gap* (unsupported) or *land* (supported) locations. A particle occupying some location in  $\Gamma$  can tell whether it is in the gap or on land. We also introduce *objects*, or static particles that do not perform computation; these are used to constrain the particles to remain connected to certain fixed sites. In order to analyze the strength of the solutions our algorithm produces, for any particle configuration  $\sigma$  and parameter  $c > 1$ , we define the *weighted perimeter*  $\bar{p}(\sigma, c)$  to be the summed edge weights of edges on the external boundary of  $\sigma$ , where edges on land have weight 1, edges in the gap have weight  $c > 1$ , and edges with one endpoint on land and one endpoint in the gap have weight  $(1+c)/2$ . Stated another way,  $\bar{p}(\sigma, c) = p(\sigma) + (c-1)g(\sigma)$ , where the *gap perimeter*  $g(\sigma)$  is the number of perimeter edges that are in the gap, where edges with one endpoint in the gap and one endpoint on land count as half an edge in the gap. Note that an edge may appear twice in a walk around the outer boundary of  $\sigma$ , and thus may be counted twice in  $\bar{p}(\sigma, c)$ ,  $p(\sigma)$  or  $g(\sigma)$ .

In the *shortcut bridging problem*, we consider an instance  $(L, O, \sigma_0, c, \alpha)$ , where  $L \subseteq V$  is the set of land locations,  $O$  is the set of (two) objects to bridge between,  $\sigma_0$  is the initial configuration of the particle system,  $c > 1$  is a fixed weight for gap edges, and  $\alpha > 1$  is a parameter capturing our error tolerance. An instance is valid if (i) the objects of  $O$  and particles of  $\sigma_0$  all occupy locations in  $L$ , (ii)  $\sigma_0$  connects the objects, and (iii)  $\sigma_0$



Figure 5.13: (a) In this image from [98], army ants of the genus *Eciton* build a dynamic bridge which balances the benefit of a shortcut path with the cost of committing ants to the structure. (b) Our shortcut bridging algorithm also balances competing objectives and converges to similar configurations.

is connected, a notion formally defined in Section 5.2.2. A (distributed) algorithm *solves* an instance  $(L, O, \sigma_0, c, \alpha)$  if, beginning from  $\sigma_0$ , it reaches and remains in (with high probability) a set of configurations  $\Sigma^*$  such that any  $\sigma \in \Sigma^*$  has weighted perimeter  $\bar{p}(\sigma, c)$  within an  $\alpha$ -factor of its minimum possible value.

The weighted perimeter balances the trade-off observed in [98] between the competing objectives of establishing a short path between the fixed endpoints while not having too many particles in the gap. We focus on gap perimeter instead of the number of particles in the gap (which is perhaps a more natural analogy to [98]) because (i) the shortcut bridges produced with this metric more closely resemble ant structures and (ii) only particles on the perimeter can move, and thus recognize the potential risk of being in the gap.

In analogy to the apparatus used in [98] (see Figure 5.13a), we are particularly interested in the special case where  $L$  forms a V-shape,  $O$  has two objects positioned at either base of  $L$ , and  $\sigma_0$  lines the interior sides of  $L$ , as in Figure 5.14a. However, our algorithm is not limited to this setting; for example, we show simulation results for an N-shaped land mass (Figure 5.14b) in Section 5.5.




Figure 5.14: Example instances  $(L, O, \sigma_0, c, \alpha)$  of the shortcut bridging problem for which we present simulation results (Section 5.5). Light (brown) nodes are land locations, large (red) nodes are occupied by objects, and black nodes are occupied by particles.

### 5.6.3 The Stochastic Approach

In the previous sections we described and analyzed the stochastic, distributed algorithm for compression in the amoebot model we introduced in [21]; here we extend that work to show that this stochastic approach is in fact more generally applicable. As we described above, the motivation underlying the use of simple Markov chains to model programmable matter comes from statistical physics. Like a spring relaxing, physical systems favor configurations that minimize energy. In the stochastic approach to programmable matter, we introduce a Hamiltonian  $H(\sigma)$  over particle configurations  $\sigma$  that assigns the lowest energy values to desirable configurations; we then design a Markov chain to favor these low energy configurations. For compression, we used the Hamiltonian  $H(\sigma) = p(\sigma)$ , the perimeter of  $\sigma$ , because our desirable configurations had short perimeter. For shortcut bridging, we use the Hamiltonian  $H(\sigma) = \bar{p}(\sigma, c)$ , the weighted perimeter of particle configuration  $\sigma$ . We then let  $w(\sigma, c) = \lambda^{-\bar{p}(\sigma, c)}$  for some parameter  $\lambda$ . As  $\lambda$  gets larger, we increasingly favor configurations where  $\bar{p}(\sigma, c)$  is small and the desired bridging behavior is exhibited.

By harnessing our knowledge of how the probabilities of local moves induce global emergent behavior, we can design a Markov chain  $\mathcal{M}$  that uses only local moves and eventually reaches a distribution that favors configurations proportional to their weight  $w(\sigma)$ . This Markov chain  $\mathcal{M}$  for *shortcut bridging* in the geometric amoebot model translates directly to a fully distributed, local, asynchronous algorithm  $\mathcal{A}$ . Rather than terminating

the process at some point and using the configuration at that time step as a random sample, we instead run the Markov chain indefinitely, moving among different configurations but remaining at the stationary distribution of  $\mathcal{M}$ , which we will prove meets our desired objectives with high probability. That is, we prove that  $\mathcal{M}$  (and by extension,  $\mathcal{A}$ ) solves the shortcut bridging problem: for any constant  $\alpha > 1$ , the long run probability that  $\mathcal{M}$  is in a configuration  $\sigma$  with  $\bar{p}(\sigma, c)$  larger than  $\alpha$  times its minimum possible value is exponentially small. More specifically, we prove (Theorem 5.40) that setting  $\lambda > 2 + \sqrt{2}$  forces the low energy configurations with small  $\bar{p}(\sigma, c)$  to dominate the state space. That is, for  $\lambda > 2 + \sqrt{2}$ , low energy configurations have sufficiently large weight that these configurations occur with high probability in the stationary distribution of  $\mathcal{M}$ . As above, the key tool used to establish this is a careful *Peierls argument*, used in statistical physics to study non-uniqueness of limiting Gibbs measures and in computer science to establish slow mixing of Markov chains.

We also specifically consider V-shaped land masses with an object on each branch of the V, and prove that the resulting bridge structures vary with the interior angle of the V-shaped gap being shortcut — a phenomenon also observed by Reid et al. [98] in the army ant bridges — and show in simulation that they are qualitatively similar to those of the ants (e.g., Figure 5.13).

As discussed in Section 5.4.3, compared to other algorithms for programmable matter and self-organizing particle systems, this stochastic method produces distributed algorithms that are nearly oblivious, more robust to failures, and require little to no communication between particles. Since these algorithms are derived from stochastic processes, powerful tools developed to analyze Markov chains can be employed to rigorously understand their behavior.

## 5.7 Algorithms for Shortcut Bridging

Recall that for the shortcut bridging problem, we want our algorithm to achieve small weighted perimeter, where boundary edges in the gap cost more than those on land. The algorithm must balance the competing objectives of having a short path between the two objects while not forming too large of a bridge. We capture these two factors by preferring small perimeter and small gap perimeter, respectively. While these objectives may appear to be aligned rather than competing, decreasing the length of the overall perimeter increases the gap perimeter and vice versa in the problem instances we consider (e.g., Figure 5.14).

Specifically, our Markov chain algorithm incorporates two bias parameters:  $\lambda$  and  $\gamma$ . The value of  $\lambda$  controls the preference for having small perimeter, while  $\gamma$  controls the preference for having small gap perimeter. In this paper, we only consider  $\lambda > 1$  and  $\gamma > 1$ , which correspond to favoring smaller perimeter and smaller gap perimeter, respectively. Using a Metropolis filter, we ensure that our algorithm converges to a distribution over particle system configurations where the relative likelihood of the particle system being in configuration  $\sigma$  is  $\lambda^{-p(\sigma)}\gamma^{-g(\sigma)}$ , or equivalently,  $\lambda^{-\bar{p}(\sigma,c)}$  for  $c = 1 + \log_\lambda \gamma$ . We note  $\lambda$  is the same parameter that controlled compression in Sections 5.3 to 5.5, where particle configurations converged to a distribution proportional to  $\lambda^{-p(\sigma)}$ . As we know that  $\lambda > 1$  is not sufficient to ensure compression, we restrict our attention to  $\lambda > 2 + \sqrt{2}$ , the regime where compression provably occurs.

To ensure our algorithm maintains some desired invariants throughout its execution, we again require that each particle move satisfies Property 1 or Property 2 (see Section ??). These properties maintain system connectivity<sup>2</sup>, prevent holes from forming, and ensure reversibility of the Markov chain; more details can be found above. These last two conditions are necessary for applying established tools from Markov chain analysis.

We can now introduce the Markov chain  $\mathcal{M}$  for an instance  $(L, O, \sigma_0, c, \alpha)$  of shortcut

---

<sup>2</sup>Since particles treat objects as static particles, the particle system may actually disconnect into several components which remain connected through objects.

bridging. For input parameter  $\lambda > 2 + \sqrt{2}$ , set  $\gamma = \lambda^{c-1}$ . Markov chain  $\mathcal{M}$  with parameters  $\lambda$  and  $\gamma$  is then as follows. We assume initial configuration  $\sigma_0$  is connected and hole-free.<sup>3</sup>

---

**Algorithm  $\mathcal{M}$ :** Markov Chain for Shortcut Bridging

---

**Beginning at any connected configuration  $\sigma_0$  of  $n$  contracted particles, repeat:**

- 1: Select particle  $P$  uniformly at random from among all particles; let  $\ell$  be its location. Choose a neighboring location  $\ell'$  and  $q \in (0, 1)$  uniformly at random.
  - 2: **if**  $\ell'$  is unoccupied **then**
  - 3:      $P$  expands to simultaneously occupy  $\ell$  and  $\ell'$ .
  - 4: **else** Return to Step 1.
  - 5: Let  $\sigma$  be the configuration with  $P$  contracted at  $\ell$  and  $\sigma'$  the configuration with  $P$  contracted at  $\ell'$ .
  - 6: **if** (1)  $|N(\ell)| \neq 5$ , (2) locations  $\ell$  and  $\ell'$  satisfy Property 1 or Property 2, and (3)  $q < \lambda^{p(\sigma)-p(\sigma')}\gamma^{g(\sigma)-g(\sigma')}$  **then**
  - 7:      $P$  contracts to  $\ell'$ .
  - 8: **else**  $P$  contracts back to  $\ell$ .
- 

Although values  $p(\sigma) - p(\sigma')$  and  $g(\sigma) - g(\sigma')$  used to check Condition (3) in Step 6 of  $\mathcal{M}$  are values defined at system-level scale, we show these differences can be calculated locally.

**Lemma 5.34.** *An expanded particle  $P$  occupying adjacent locations  $\ell$  and  $\ell'$  in  $\Gamma$  can calculate the values of  $p(\sigma) - p(\sigma')$  and  $g(\sigma) - g(\sigma')$  in Step 3(iii) of  $\mathcal{M}$  using only local information involving  $\ell$ ,  $\ell'$ , and  $N(\ell \cup \ell')$ .*

*Proof.* Observe that these values need only be calculated if Conditions (1) and (2) in Step 6 of  $\mathcal{M}$  hold. By our work above on compression,

$$p(\sigma) - p(\sigma') = |N(\ell')| - |N(\ell)|,$$

which can be calculated using only local information.

Recall that gap perimeter is defined as the number of perimeter edges in the gap, counting edges between gap and land as half an edge; this is equal to the number of particles

---

<sup>3</sup>If  $\sigma_0$  has holes, our algorithm will eliminate them and they will not reform, as discussed at length for our compression algorithm above; here we focus only on the behavior of the system after this occurs.

that are on the perimeter and in the gap, counted with appropriate multiplicity if a particle appears on the perimeter more than once. Given a particle  $Q$  and a configuration  $\tau$ , let  $G(Q, \tau)$  be equal to 1 if  $Q$  occupies a gap location in  $\tau$  and 0 otherwise. Let  $\delta(Q, \tau)$  be the number of times  $Q$  appears on the perimeter of  $\tau$ . Then the desired difference is:

$$g(\sigma) - g(\sigma') = \sum_Q [G(Q, \sigma)\delta(Q, \sigma) - G(Q, \sigma')\delta(Q, \sigma')].$$

Define  $\Delta(Q) = \delta(Q, \sigma) - \delta(Q, \sigma')$ . For particle  $P$ , since Conditions (1) and (2) in Step 6 hold,  $\Delta(P) = 0$ . For any particle  $Q \notin \{P\} \cup N(\ell \cup \ell')$ ,  $\Delta(Q) = 0$  since its neighborhood is not affected by the movement of  $P$ . Moreover, for any particle  $Q \neq P$ ,  $G(Q, \sigma) = G(Q, \sigma')$  since it does not move. So:

$$g(\sigma) - g(\sigma') = \delta(P, \sigma) [G(P, \sigma) - G(P, \sigma')] + \sum_{Q \in N(\ell \cup \ell')} G(Q, \sigma)\Delta(Q).$$

The first term is easily calculated locally. For the summation, it remains to show that  $P$  can locally calculate  $\Delta(Q)$  for any  $Q \in N(\ell \cup \ell')$ . First suppose that  $Q$  occupies a location adjacent to  $\ell$  but not  $\ell'$ . Then:

$$\Delta(Q) = \begin{cases} -1 & \text{if } Q \text{ has two neighbors in } N(\ell), \\ 1 & \text{if } Q \text{ has no neighbors in } N(\ell), \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

The opposite is true if  $Q$  occupies a location adjacent to  $\ell'$  but not  $\ell$ . Lastly, suppose  $Q$  occupies a location adjacent to both  $\ell$  and  $\ell'$ . Then:

$$\Delta(Q) = \begin{cases} 0 & \text{if } Q \text{ has zero or two neighbors in } N(\ell \cup \ell'), \\ -1 & \text{if } Q \text{ shares a neighbor with } \ell \text{ but not } \ell', \text{ and} \\ 1 & \text{if } Q \text{ shares a neighbor with } \ell' \text{ but not } \ell. \end{cases}$$

In all cases,  $P$  can calculate  $\Delta(Q)$ , and thus also  $g(\sigma) - g(\sigma')$ , using only local information. □

The state space  $\Omega$  of  $\mathcal{M}$  is the set of all configurations reachable from  $\sigma_0$  via valid transitions of  $\mathcal{M}$ . We conjecture that when  $\sigma_0$  is a connected hole-free configuration that this includes all connected, hole-free configurations of  $n$  particles connected to both objects, but proving all such configurations are reachable from  $\sigma_0$  is not necessary for our results. The proof of the corresponding result for compression (Section 5.4.5) does not generalize due to the presence of static objects.

While  $\mathcal{M}$  is a Markov chain with centralized control of the particle system, using Lemma 5.34 one can transform  $\mathcal{M}$  into a distributed, local, asynchronous algorithm  $\mathcal{A}$  that each particle runs individually. The full details of this construction are identical to those given in Section 5.4.2 for compression, so we do not repeat them here.

### 5.7.1 Properties of Markov Chain $\mathcal{M}$

We now show some useful properties of the Markov chain  $\mathcal{M}$ . Our first two claims follow from work on the compression problem above, and basic properties of Markov chains and our particle systems.

**Lemma 5.35.** *If  $\sigma_0$  is connected and has no holes, then at every iteration of  $\mathcal{M}$ , the current configuration is connected and has no holes.*

*Proof.* For compression, we proved that no valid moves could introduce holes or disconnect the particle system; this is Lemmas 5.6 and 5.7. Since the moves allowed by our shortcut bridging algorithm  $\mathcal{M}$  are a subset of those allowed in the compression algorithm (since the local properties checked at each iteration are the same),  $\mathcal{M}$  cannot introduce holes or disconnect the system. □

**Lemma 5.36.** *For any  $\sigma$  and  $\tau$  in  $\Omega$ ,  $P(\sigma, \tau) > 0$  if and only if  $P(\tau, \sigma) > 0$ .*



*Proof.* The proof is nearly identical to that of Lemma 5.14 for compression, above. Because  $\sigma$  and  $\tau$  are both hole-free configurations, Properties 1 and 2 ensure that particle  $P$  moving from location  $\ell$  to location  $\ell'$  is valid if and only if  $P$  moving from  $\ell'$  to  $\ell$  is.  $\square$

**Lemma 5.37.** *If  $\sigma_0$  has no holes, then  $\mathcal{M}$  is ergodic.*

*Proof.* We defined  $\Omega$  to be precisely those configurations reachable by valid transitions of  $\mathcal{M}$  starting from  $\sigma_0$ . By Lemma 5.36, from every state the initial state  $\sigma_0$  is reachable. Altogether, this implies  $\mathcal{M}$  is irreducible on its state space  $\Omega$ .  $\mathcal{M}$  is aperiodic because at each iteration there is a probability of at least  $1/6$  that no move occurs, as each particle has at least one neighbor. Thus, the chain  $\mathcal{M}$  is ergodic.  $\square$

As  $\mathcal{M}$  is finite and ergodic, it converges to a unique stationary distribution, and we can find that distribution using detailed balance.

**Lemma 5.38.** *The stationary distribution of  $\mathcal{M}$  is*

$$\pi(\sigma) = \frac{\lambda^{-p(\sigma)}\gamma^{-g(\sigma)}}{Z},$$

where  $Z = \sum_{\sigma \in \Omega} \lambda^{-p(\sigma)}\gamma^{-g(\sigma)}$ .

*Proof.* Using Lemma 5.36, we verify the claim via detailed balance. Let  $\sigma, \tau \in \Omega$  be distinct configurations that differ by one valid move of a particle  $P$  from location  $\ell$  to neighboring location  $\ell'$ , and let  $n$  be the number of particles in the system. Then,

$$P(\sigma, \tau) = \frac{1}{n} \cdot \frac{1}{6} \cdot \min\{\lambda^{p(\sigma)-p(\tau)}\gamma^{g(\sigma)-g(\tau)}, 1\}, \text{ and}$$

$$P(\tau, \sigma) = \frac{1}{n} \cdot \frac{1}{6} \cdot \min\{\lambda^{p(\tau)-p(\sigma)}\gamma^{g(\tau)-g(\sigma)}, 1\}.$$

If  $\lambda$  and  $\gamma$  satisfy  $\lambda^{p(\sigma)-p(\tau)}\gamma^{g(\sigma)-g(\tau)} \leq 1$ , then

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{-p(\sigma)}\gamma^{-g(\sigma)}}{Z} \cdot \frac{\lambda^{p(\sigma)-p(\tau)}\gamma^{g(\sigma)-g(\tau)}}{6n} = \frac{\lambda^{-p(\tau)}\gamma^{-g(\tau)}}{Z} \cdot \frac{1}{6n} = \pi(\tau)P(\tau, \sigma).$$

Otherwise, it must hold that  $\lambda^{p(\sigma)-p(\tau)}\gamma^{g(\sigma)-g(\tau)} > 1$ , and we see that

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{-p(\sigma)}\gamma^{-g(\sigma)}}{Z} \cdot \frac{1}{6n} = \frac{\lambda^{-p(\tau)}\gamma^{-g(\tau)}}{Z} \cdot \frac{\lambda^{p(\tau)-p(\sigma)}\gamma^{g(\tau)-g(\sigma)}}{6n} = \pi(\tau)P(\tau, \sigma).$$

In both cases detailed balance is satisfied. Using definition of  $Z$ , we see that  $\pi$  satisfies  $\sum_{\sigma \in \Omega} \pi(\sigma) = 1$ , so  $\pi$  is a valid probability distribution and we conclude  $\pi$  is the unique stationary distribution of  $\mathcal{M}$ .  $\square$

We can also express this stationary distribution using weighted perimeter.

**Lemma 5.39.** *For  $c = 1 + \log_\lambda \gamma$ , the stationary distribution of  $\mathcal{M}$  is given by*

$$\pi(\sigma) = \lambda^{-\bar{p}(\sigma, c)} / Z,$$

where  $Z = \sum_{\sigma \in \Omega} \lambda^{-\bar{p}(\sigma, c)}$ .

*Proof.* This follows from the definition of  $\bar{p}(\sigma, c)$ .  $\square$

**Theorem 5.40.** *Consider an execution of Markov chain  $\mathcal{M}$  on state space  $\Omega$ , starting at configuration  $\sigma_0$  with  $n$  particles, where  $\lambda > 2 + \sqrt{2} =: \bar{c}$  and  $\gamma > 1$ . For any constant  $\alpha$  satisfying*

$$\alpha > \frac{\log \lambda}{\log \lambda - \log \bar{c}} > 1,$$

*the probability that particle configuration  $\sigma$  drawn at random from  $\mathcal{M}$ 's stationary distribution  $\pi$  satisfies  $\bar{p}(\sigma, 1 + \log_\lambda \gamma) > \alpha \cdot \bar{p}_{\min}$  is exponentially small in  $n$ , where  $\bar{p}_{\min}$  is the minimum weighted perimeter of a configuration in  $\Omega$ .*

*Proof.* This proof mimics that of  $\alpha$ -compression in [21], but additional insights and care were necessary to accommodate the difficulties introduced by considering weighted perimeter. Throughout we consider weighted perimeter  $\bar{p}(\sigma) = \bar{p}(\sigma, 1 + \log_\lambda \gamma)$ .

Define the weight of a configuration  $\sigma \in \Omega$  to be:

$$w(\sigma) := \pi(\sigma) \cdot Z = \lambda^{-p(\sigma)}\gamma^{-g(\sigma)} = \lambda^{-\bar{p}(\sigma)}.$$

For a set of configurations  $S \subseteq \Omega$ , we define its weight  $w(S) = \sum_{\sigma \in S} w(\sigma)$ . Let  $\sigma_{min} \in \Omega$  be a configuration of  $n$  particles with minimal weighted perimeter  $\bar{p}_{min}$ , and let  $S_\alpha$  be the set of configurations with weighted perimeter at least  $\alpha \cdot \bar{p}_{min}$ . We show:

$$\pi(S_\alpha) = \frac{w(S_\alpha)}{Z} < \frac{w(S_\alpha)}{w(\sigma_{min})} \leq f_2(n)\zeta^{\sqrt{n}},$$

where  $\zeta < 1$  and  $f_2(n)$  is a function that is subexponential in  $n$ . The first equality and inequality follow directly from the definitions of  $Z$ ,  $w$ , and  $\sigma_{min}$ . We focus on the last inequality.

Stratify  $S_\alpha$  into sets of configurations that have the same weighted perimeter; there are at most  $\mathcal{O}(n^2)$  such sets, as the total perimeter and gap perimeter can each take on at most  $\mathcal{O}(n)$  integer values. Label these sets as  $A_1, A_2, \dots, A_m$  in order of increasing weighted perimeter, where  $m$  is the total number of distinct weighted perimeters of configurations in  $S_\alpha$ . Let  $\bar{p}_i$  be the weighted perimeter of all configurations in set  $A_i$ ; since  $A_i \subseteq S_\alpha$ , then  $\bar{p}_i \geq \alpha \cdot \bar{p}_{min}$ .

Note  $w(\sigma) = \lambda^{-\bar{p}_i}$  for every  $\sigma \in A_i$ , so to bound  $w(A_i)$  it suffices to bound  $|A_i|$ . A configuration with weighted perimeter  $\bar{p}_i$  has perimeter  $p \leq \bar{p}_i$ , and Lemma 5.23 implies the number of connected, hole-free particle configurations with perimeter  $p$  is at most  $f(p)\bar{c}^p$ , for some subexponential function  $f$ . Letting  $p_{min}$  denote the minimum possible (unweighted) perimeter of a configuration of  $n$  particles, we conclude that:

$$w(A_i) = \lambda^{-\bar{p}_i} |A_i| \leq \lambda^{-\bar{p}_i} \cdot \sum_{p=p_{min}}^{\bar{p}_i} f(p)\bar{c}^p \leq \lambda^{-\bar{p}_i} f_1(\bar{p}_i)\bar{c}^{\bar{p}_i},$$

where  $f_1(\bar{p}_i) = \sum_{p=p_{min}}^{\bar{p}_i} f(p)$  is necessarily also a subexponential function because it is a sum of at most a linear number of subexponential terms. So,

$$w(S_\alpha) = \sum_{i=1}^m w(A_i) \leq \sum_{i=1}^m f_1(\bar{p}_i) \left(\frac{\bar{c}}{\lambda}\right)^{\bar{p}_i} \leq f_2(n) \left(\frac{\bar{c}}{\lambda}\right)^{\alpha \bar{p}_{min}},$$

where  $f_2(n) = \sum_{i=1}^m f_1(\bar{p}_i)$  is a subexponential function because  $\bar{p}_i = \mathcal{O}(n)$ ,  $m = \mathcal{O}(n^2)$ , and  $f_1$  is subexponential. The last inequality above holds as  $\lambda > \bar{c}$  and  $\bar{p}_i \geq \alpha \cdot \bar{p}_{min}$  because  $A_i \subseteq S_\alpha$ . Then, since  $w(\sigma_{min}) = \lambda^{-\bar{p}_{min}}$ ,

$$\pi(S_\alpha) < \frac{w(S_\alpha)}{w(\sigma_{min})} \leq f_2(n) \left(\frac{\bar{c}}{\lambda}\right)^{\alpha \bar{p}_{min}} \lambda^{\bar{p}_{min}} = f_2(n) \left[\lambda \left(\frac{\bar{c}}{\lambda}\right)^\alpha\right]^{\bar{p}_{min}}.$$

The constant  $\zeta = \lambda(\bar{c}/\lambda)^\alpha$  is less than one whenever  $\alpha > \frac{\log \lambda}{\log \lambda - \log \bar{c}}$ . We have  $\bar{p}_{min} \geq \sqrt{n}$  because any  $n$  particles must have perimeter at least  $\sqrt{n}$  by Lemma 5.1. This suffices to show there is a constant  $\zeta < 1$  and a subexponential function  $f_2(n)$  such that  $\pi(S_\alpha) < f_2(n)\zeta^{\sqrt{n}}$ , which proves the theorem.  $\square$

The following corollary shows that our algorithm solves any instance  $(L, O, \sigma_0, c, \alpha)$  of the shortcut bridging problem when parameters  $\lambda$  and  $\gamma$  are chosen accordingly.

**Corollary 5.41.** *The distributed, local algorithm  $\mathcal{A}$  associated with Markov chain  $\mathcal{M}$  solves any valid instance of the shortcut bridging problem.*

*Proof.* Given any valid instance  $(L, O, \sigma_0, c, \alpha)$  of the shortcut bridging problem, it suffices to run  $\mathcal{A}$  starting from configuration  $\sigma_0$  with parameters  $\lambda > (2 + \sqrt{2})^{\frac{\alpha}{\alpha-1}}$  and  $\gamma = \lambda^{c-1}$ . Then  $\alpha > \frac{\log(\lambda)}{\log(\lambda) - \log(2 + \sqrt{2})} > 1$ , so by Theorem 5.40 the system reaches and remains with all but exponentially small probability in a set of configurations with weighted perimeter  $\bar{p}(\sigma, c) \leq \alpha \cdot \bar{p}_{min}$ , where  $\bar{p}_{min}$  is the minimum weighted perimeter of a configuration in  $\Omega$ .  $\square$

### 5.7.2 Simulations

We can see the performance of our algorithm from simulation results on a variety of instances. Figure 5.15 shows snapshots over time for a bridge shortcutting a V-shaped gap with internal angle  $\theta = \pi/3$  and biases  $\lambda = 4, \gamma = 2$ . Qualitatively, this bridge matches the shape and position of the army ant bridges in [98]. Figure 5.16 shows the resulting bridge structure when the land mass is N-shaped. Lastly, Figure 5.17 shows the results of

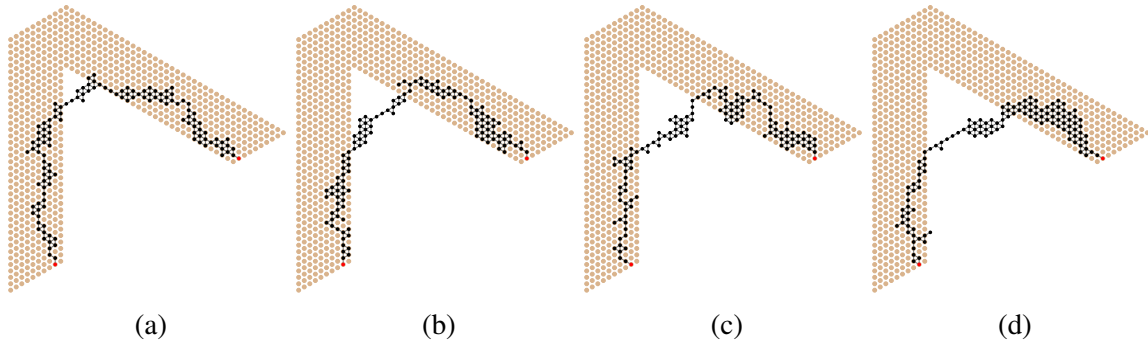


Figure 5.15: A particle system using biases  $\lambda = 4$  and  $\gamma = 2$  to shortcut a V-shaped land mass with  $\theta = \pi/3$  after (a) 2 million, (b) 4 million, (c) 6 million, and (d) 8 million iterations of Markov chain  $\mathcal{M}$ , beginning in configuration  $\sigma_0$  shown in Figure 5.14a.




Figure 5.16: A particle system using biases  $\lambda = 4$  and  $\gamma = 2$  to shortcut an N-shaped land mass after (a) 10 million and (b) 20 million iterations of Markov chain  $\mathcal{M}$ , beginning in configuration  $\sigma_0$  shown Figure 5.14b.

an experiment that held  $\lambda$ ,  $\gamma$ , and the number of iterations of  $\mathcal{M}$  constant, varying only the internal angle of the V-shaped land mass. The particle system exhibits behavior consistent with the theoretical results to come in Section 5.8 and the army ant bridges in [98], short-cutting closer to the bottom of the gap when  $\theta$  is small and staying almost entirely on land when  $\theta$  is large.

These simulations demonstrate the successful application of our stochastic approach to shortcut bridging. Moreover, experimenting with variants suggests this approach may be useful for other related applications in the future.




Figure 5.17: A particle system using biases  $\lambda = 4$  and  $\gamma = 2$  to shortcut a V-shaped land mass with angle (a)  $\pi/6$ , (b)  $\pi/3$ , and (c)  $\pi/2$  after 20 million iterations of Markov chain  $\mathcal{M}$ . For a given angle, the land mass  $L$  and initial configuration  $\sigma_0$  were constructed as described in Section 5.8.

### 5.8 Dependence of Bridge Structure On Gap Angle

To understand the relationship between bridging and shape, we consider V-shaped land masses of various angles (e.g., Figure 5.14a). We prove our shortcut bridging algorithm has a dependence on the internal angle  $\theta$  of the gap similar to that of the army ant bridges studied by Reid et al. [98]. We show that when  $\theta$  is sufficiently small, with all but exponentially small probability the bridge constructed by the particles stays close to the bottom of the gap (away from the apex of angle  $\theta$ ). On the other hand, we show that for some large values of  $\theta$ , when  $\lambda$  and  $\gamma$  satisfy certain conditions, with all but exponentially small probability the bridge stays close to the top of the gap. We prove these results with a Peierls argument and careful analysis of the geometry of the gap. Throughout this section we will measure angles in radians, rather than degrees.

We first formalize our V-shaped land mass  $L$  for any  $\theta \in (0, \pi)$ ; see Figure 5.18a ( $\theta \sim \pi/6$ ) and Figure 5.18b ( $\theta \sim \pi/2$ ). Let  $e \in E$  be any edge of the triangular lattice and label its endpoints as  $v_1$  and  $v_2$ . Extend line segment  $\ell_1$  from  $v_1$  such that it forms an angle of  $\pi/2 + \theta/2$  with  $e$ . Similarly extend line segment  $\ell_2$  from  $v_2$ , of the same length and on the same side of  $e$  as  $\ell_1$ , also forming an angle of  $\pi/2 + \theta/2$  with  $e$ . Segments  $\ell_1$  and  $\ell_2$  then differ in their orientation by angle  $\theta$ . Without loss of generality, we assume  $\ell_1$  is clockwise from  $\ell_2$  around  $e$ . The land mass consists of  $v_1, v_2$ , and all vertices of  $\Gamma$  outside

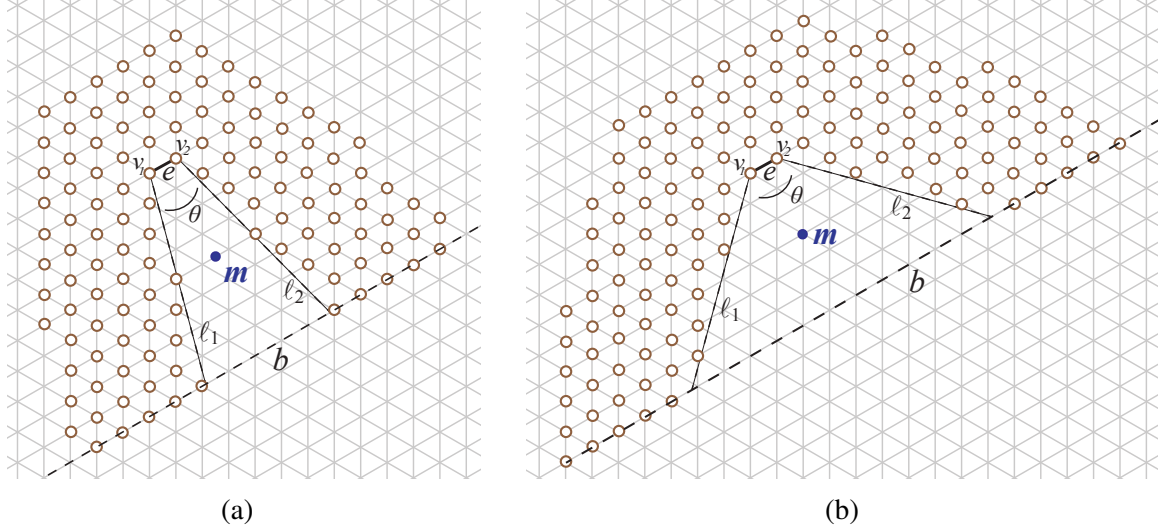


Figure 5.18: The land mass  $L$  of constant width 5 for (a) a small value of  $\theta \sim \pi/6$  and height 8 and (b) a large value of  $\theta \sim \pi/2$  and height 9. Point  $m$  is the midpoint of the segment between the midpoints of  $\ell_1$  and  $\ell_2$ , and  $b$  is shown as a dashed line.

of  $\ell_1$  and  $\ell_2$  up to some constant width; e.g., in Figure 5.18 that width is five. This careful definition involving edge  $e$  is necessary to ensure there are no adjacent land locations on opposite sides of the gap, as could happen for small  $\theta$  if the land mass is not constructed carefully.

From now on we will, in a slight abuse of notation, refer to the gap locations between  $\ell_1$  and  $\ell_2$  as *the gap*. By the *bottom of the gap*, we mean the line  $b$  through  $\ell_1$  and  $\ell_2$ 's other endpoints (not  $v_1$  and  $v_2$ ). We may assume  $b$  is a line of the triangular lattice by truncating  $\ell_1$  and  $\ell_2$  so that both end on a lattice line; this does not change the land mass  $L$ . We also assume  $b \cap \ell_1$  and  $b \cap \ell_2$  are not vertices of the triangular lattice  $\Gamma$ ; if they are, we can perturb  $\ell_1$  and  $\ell_2$  slightly, without changing the land mass. Note  $b$  is always parallel to  $e$ .

The *height* of land mass  $L$  is the length of a shortest path in  $\Gamma$  from  $v_1$  or  $v_2$  to  $b$  that only visits land locations; the land mass in Figure 5.18a has height 8, while the land mass in Figure 5.18b has height 9. Let  $m$  be the midpoint of the segment connecting the midpoints of  $\ell_1$  and  $\ell_2$ ;  $m$  is in the center of the gap, halfway between  $e$  and  $b$ .

The initial configuration  $\sigma_0$  we consider is a path of width 2 lining the interior sides of the land mass  $L$ ; see Figure 5.19. We position the two fixed objects of  $O$  in line  $b$  at the

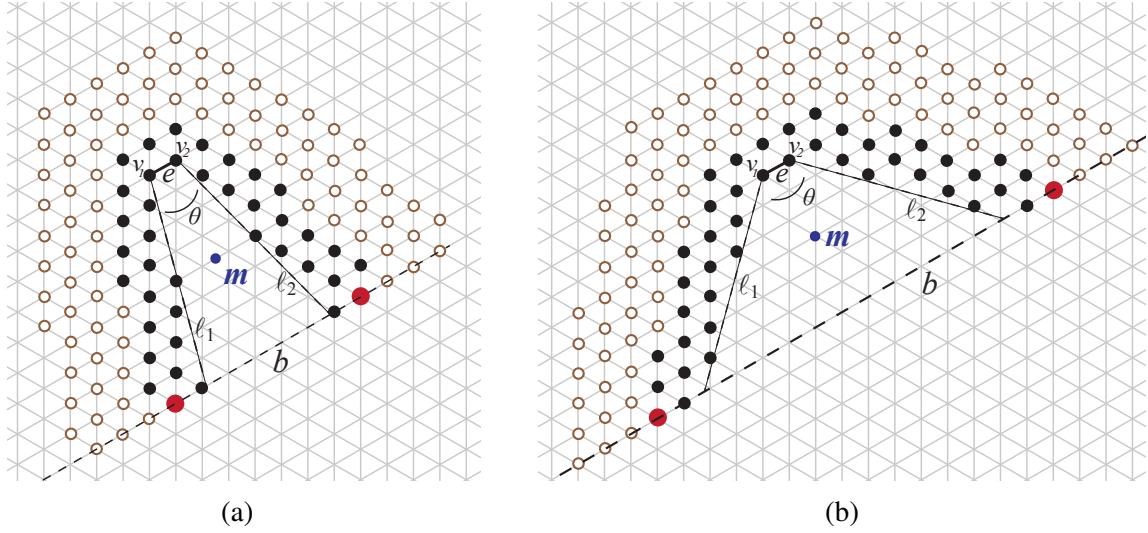


Figure 5.19: The initial configuration  $\sigma_0$ , with particles shown in black and objects enlarged and red, for (a) a small value of  $\theta \sim \pi/6$  and (b) a large value of  $\theta \sim \pi/2$ . Point  $m$  is the midpoint of the segment between the midpoints of  $\ell_1$  and  $\ell_2$ , and  $b$  is shown as a dashed line.

second vertices outside  $\ell_1$  and  $\ell_2$ , anchoring the particles on either side of the gap. Note the height of  $L$  is exactly the number of particles in  $\sigma_0$  next to  $\ell_1$  (or  $\ell_2$ ), excluding  $v_1$  and  $v_2$ .

**Lemma 5.42.** *Let  $L$  be a V-shaped land mass of height  $k$  and angle  $\theta$ . The initial configuration  $\sigma_0$  has  $4k + 5$  particles and two objects.*

*Proof.* First, suppose  $\theta \leq \pi/3$ , as in Figure 5.19a. Each lattice line parallel to  $e$  and intersecting  $\ell_1$  and  $\ell_2$ , up to but not including  $b$ , contains exactly four particles. There are  $k$  such lattice lines. Line  $b$  contains two particles. In the lattice line above and parallel to  $e$ , there are three particles. In total, this gives  $4k + 2 + 3 = 4k + 5$  particles and two objects.

Now, suppose  $\theta > \pi/3$ , as in Figure 5.19b; a different counting approach is necessary. Consider the lattice line through  $v_1$  and the gap location adjacent to  $v_1$  and  $v_2$ ; this line and all lines parallel to it intersecting  $\ell_1$  contain exactly two particles, and there are  $k$  such lines. The same is true for  $v_2$  and  $\ell_2$ . Uncounted by this approach are five additional particles: the two particles adjacent to each of the two objects, and the particle adjacent to  $v_1$  and  $v_2$ . In total, this gives  $2k + 2k + 4 + 1 = 4k + 5$  particles and two objects.  $\square$



For a given  $\sigma$ , let  $x$  be the particle or object contained in line  $b$  farthest outside of  $\ell_1$ , and let  $y$  be the particle or object in line  $b$  farthest outside of  $\ell_2$ . We will refer to the perimeter of  $\sigma$  traversed counterclockwise from  $x$  to  $y$  as the *inner perimeter* of  $\sigma$ , and the perimeter of  $\sigma$  traversed clockwise from  $x$  to  $y$  as the *outer perimeter* of  $\sigma$ . We say the inner perimeter is *above a point*  $p$  if  $p$  is to the right of the inner perimeter traversed from  $x$  to  $y$ ; it is *below a point*  $p$  if  $p$  is to its left.

We can partition  $\Omega$  into two sets  $S_1$  and  $S_2$ , where  $S_1$  contains all configurations whose inner perimeter is strictly above midpoint  $m$  of the gap and  $S_2$  contains all configurations whose inner perimeter goes through or below  $m$ . We first prove that for  $\lambda > 2 + \sqrt{2}$  (i.e., in the range of compression) and  $\gamma > 1$ , there is an angle  $\theta_1$  such that for all  $\theta < \theta_1$ ,  $\pi(S_1)$  is exponentially small. We then prove that for  $\lambda > 2 + \sqrt{2}$  and  $\gamma > \lambda^4(2 + \sqrt{2})^4$ , there is a  $\theta_2$  such that for all  $\theta \in (\pi/3, \theta_2)$ ,  $\pi(S_2)$  is exponentially small. We expect much better bounds  $\theta_1$  and  $\theta_2$  can be obtained with more effort, and that these results generalize to all  $\lambda > 2 + \sqrt{2}$  and  $\gamma > 1$ , but here we simply demonstrate it is possible to give rigorous results about the dependence of the bridge structure on  $\theta$ .

### 5.8.1 Proofs for Small $\theta$

We begin with some structural lemmas.

**Lemma 5.43.** *Let  $L$  be a V-shaped land mass of height  $k$  and angle  $\theta \leq \pi/3$ . Then any path in  $\Gamma$  that starts and ends at the bottom of gap and goes strictly above the midpoint  $m$  of the gap has length at least  $k + 1$ .*

*Proof.* For  $\theta \leq \pi/3$ , there are  $k - 1$  lattice lines parallel to  $b$  strictly between  $b$  and  $e$ . Of these lines exactly  $\lceil (k - 1)/2 \rceil$  are below or contain  $m$ . Any path from  $b$  to a location above  $m$  and back to  $b$  must contain at least two vertices in each of these lattice lines, two vertices in  $b$ , and one vertex strictly above  $m$ , giving a total of  $3 + 2\lceil (k - 1)/2 \rceil \geq 3 + 2((k - 1)/2) = k + 2$  vertices. As the length of a path is the number of edges it contains, the path must have length at least  $k + 1$ . □

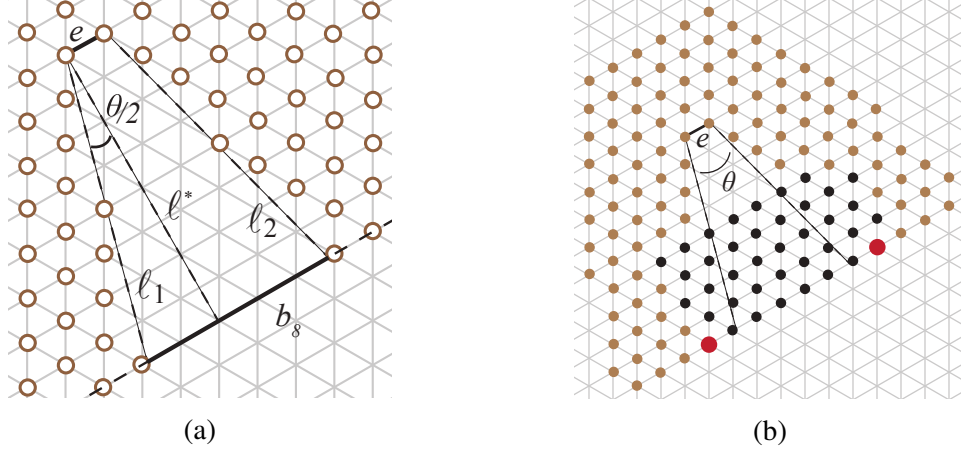


Figure 5.20: Figures from proofs in Section 5.8.1. (a) A depiction of the notation used in the proof of Lemma 5.44; the intersection of  $b_8$  and the gap is depicted as a solid segment, which is of length  $8\sqrt{3}\tan(\theta/2) + 1$  and contains 4 gap locations. (b) The configuration  $\sigma^*$  used in Lemma 5.45 for  $\theta = \pi/6$  and  $k = 8$ .

**Lemma 5.44.** *The  $i$ -th lattice line below and parallel to  $e$  contains  $h(i)$  gap locations between  $\ell_1$  and  $\ell_2$ , where*

$$i\sqrt{3}\tan\frac{\theta}{2} \leq h(i) \leq i\sqrt{3}\tan\frac{\theta}{2} + 2.$$

*Proof.* Let  $b_i$  be the  $i$ -th lattice line below and parallel to  $e$ . We use trigonometry to analyze the length of  $b_i$  between  $\ell_1$  and  $\ell_2$ ; see Figure 5.20a. Consider the triangle formed by  $b_i$ ,  $\ell_1$ , and the line perpendicular to  $e$  at  $v_1$ , which we call  $\ell^*$ . Lines  $\ell_1$  and  $\ell^*$  form an angle of  $\theta/2$ , and the distance between  $e$  and  $b_i$  along  $\ell^*$  is  $i\sqrt{3}/2$ . It follows that the length of  $b_i$  between  $\ell_1$  and  $\ell^*$  is  $i\sqrt{3}\tan(\theta/2)/2$ . Altogether, this implies  $b_i$  between  $\ell_1$  and  $\ell_2$  is of length  $i\sqrt{3}\tan(\theta/2) + 1$ . As each edge of the triangular lattice is length 1, this means there are between  $i\sqrt{3}\tan(\theta/2)$  and  $i\sqrt{3}\tan(\theta/2) + 2$  gap locations in  $b_i$ , as claimed.  $\square$

**Lemma 5.45.** *Let  $L$  be a V-shaped land mass of height  $k$  and angle  $\theta \leq \pi/3$ . Then the normalizing constant  $Z$  of the stationary distribution  $\pi$  of  $\mathcal{M}$  satisfies*

$$Z \geq C \left[ (\lambda\gamma)^{-2\sqrt{3}\tan\frac{\theta}{2}} \right]^k,$$

for a constant  $C$  that depends on  $\theta$ ,  $\lambda$ , and  $\gamma$  but not on  $k$ .

*Proof.* Observe that  $Z = \sum_{\sigma \in \Omega} \lambda^{-p(\sigma)} \gamma^{-g(\sigma)}$  satisfies  $Z \geq \lambda^{-p(\sigma^*)} \gamma^{-g(\sigma^*)}$  for any  $\sigma^* \in \Omega$ . We now construct a particular  $\sigma^*$  (Figure 5.20b) and calculate its perimeter and gap perimeter. Let  $\sigma^*$  contain a straight line of particles along  $b$  connecting the two objects, and let  $u$  be the number of objects and particles in this line. By Lemma 5.44, since  $b = b_k$  and  $u$  includes two particles on land as well as two objects,

$$k\sqrt{3} \tan \frac{\theta}{2} + 4 \leq u \leq k\sqrt{3} \tan \frac{\theta}{2} + 6.$$

Continue constructing  $\sigma^*$  by placing rows of  $u$  particles above this initial row such that the row starts and ends on opposite sides of the gap. By Lemma 5.42, there are  $4k + 7$  total objects and particles, so there will be  $v = \lceil (4k + 7)/u \rceil$  such rows, with the last row possibly incomplete. We note that  $v$  satisfies:

$$\begin{aligned} v = \left\lceil \frac{4k + 7}{u} \right\rceil &\leq \frac{4k + 7}{u} + 1 \leq \frac{4k + 7}{k\sqrt{3} \tan \frac{\theta}{2} + 4} + 1 \leq \frac{4}{\sqrt{3} \tan \frac{\theta}{2}} + \frac{7}{4} + 1 \leq \frac{4}{\sqrt{3} \tan \frac{\theta}{2}} + 3 \\ v = \left\lceil \frac{4k + 7}{u} \right\rceil &\geq \frac{4k + 7}{u} \geq \frac{4k + 7}{k\sqrt{3} \tan \frac{\theta}{2} + 6} \geq \frac{4k}{k\sqrt{3} \tan \frac{\theta}{2} + 6k} \geq \frac{4}{\sqrt{3} \tan \frac{\theta}{2} + 6} \end{aligned}$$

Configuration  $\sigma^*$  has perimeter at most  $2u + 2v - 4$  and gap perimeter at most  $u - 4 + z$ , where  $z$  is the number of particles occupying gap locations in the upper perimeter of  $\sigma^*$ . These  $z$  remaining particles must be in either the  $(k - v + 1)$ -th or  $(k - v + 2)$ -th lattice lines below  $e$ , so we can bound  $z$  by again applying Lemma 5.44:

$$z \leq (k - v + 1)\sqrt{3} \tan \frac{\theta}{2} + 2.$$

Altogether, this implies:

$$\begin{aligned}
p(\sigma^*) &\leq 2u + 2v - 4 \\
&\leq 2k\sqrt{3}\tan\frac{\theta}{2} + 12 + \frac{8}{\sqrt{3}\tan\frac{\theta}{2}} + 6 - 4 \\
&\leq k\left(2\sqrt{3}\tan\frac{\theta}{2}\right) + \left(\frac{8}{\sqrt{3}\tan\frac{\theta}{2}} + 14\right),
\end{aligned}$$

and

$$\begin{aligned}
g(\sigma^*) &\leq u - 4 + z \\
&\leq k\sqrt{3}\tan\frac{\theta}{2} + 6 - 4 + (k - v + 1)\sqrt{3}\tan\frac{\theta}{2} + 2 \\
&\leq 2k\sqrt{3}\tan\frac{\theta}{2} + \left(-\frac{4}{\sqrt{3}\tan\frac{\theta}{2} + 6} + 1\right)\sqrt{3}\tan\frac{\theta}{2} + 4 \\
&\leq k\left(2\sqrt{3}\tan\frac{\theta}{2}\right) + \left(\sqrt{3}\tan\frac{\theta}{2} - \frac{4\sqrt{3}\tan\frac{\theta}{2}}{\sqrt{3}\tan\frac{\theta}{2} + 6} + 4\right).
\end{aligned}$$

We note that the second parentheses in the final bounds above for  $p(\sigma^*)$  and  $g(\sigma^*)$  are constants that only depend on  $\theta$ . This implies that there is a constant

$$C = \lambda^{-\left(14 + \frac{8}{\sqrt{3}\tan\frac{\theta}{2}}\right)} \gamma^{-\left(\sqrt{3}\tan\frac{\theta}{2} - \frac{4\sqrt{3}\tan\frac{\theta}{2}}{\sqrt{3}\tan\frac{\theta}{2} + 6} + 4\right)}$$

such that

$$Z \geq \lambda^{-p(\sigma^*)} \gamma^{-g(\sigma^*)} \geq C \left[ (\lambda\gamma)^{-2\sqrt{3}\tan\frac{\theta}{2}} \right]^k.$$

As claimed,  $C$  depends only on  $\lambda$ ,  $\gamma$ , and  $\theta$ , and is independent of  $k$ . □

**Theorem 5.46.** *Let  $\lambda > 2 + \sqrt{2} =: \bar{c}$  and  $\gamma > 1$ . Then there exists a constant  $\theta_1$  such that for all V-shaped land masses with angle  $\theta < \theta_1$ , the probability that the inner perimeter is above midpoint  $m$  is exponentially small in  $k$ , the height of the gap. In particular,*

$$\theta_1 = 2 \tan^{-1} \left( \frac{\log_{\lambda\gamma}(\lambda/\bar{c})}{\sqrt{3}} \right).$$

*Proof.* Recall that  $S_1 \subseteq \Omega$  is the set of configurations for which the inner perimeter is strictly above  $m$ . We show that  $S_1$  has exponentially small weight at stationarity; in particular, we show  $\pi(S_1)$  is bounded above by  $f_2(k)\xi^k$ , where  $f_2(k)$  is a subexponential function and  $\xi < 1$  is a constant.

If  $\sigma \in S_1$ , then by Lemma 5.43 we have  $p(\sigma) \geq 2k + 2$ , as both its inner and outer perimeters must go above  $m$ . Furthermore, because the perimeter by definition includes both objects and particles, which number  $4k + 7$  by Lemma 5.42, any configuration  $\sigma \in \Omega$  has  $p(\sigma) \leq 2(4k + 7) - 2 = 8k + 12$ . Lemma 5.23 shows the number of connected, hole-free particle configurations with perimeter  $p$  is at most  $f(p)(\bar{c})^p$  for some subexponential function  $f$ . This is certainly also an upper bound on the number of configurations in  $S_1$  with perimeter  $p$ . Because  $\gamma^{-g(\sigma)} < 1$ , we have:

$$\pi(S_1) = \sum_{\sigma \in S_1} \frac{\lambda^{-p(\sigma)} \gamma^{-g(\sigma)}}{Z} < \sum_{p=2k+2}^{8k+12} \frac{f(p) \bar{c}^p \lambda^{-p}}{Z}.$$

Let  $f_1(k) = \sum_{p=2k+2}^{8k+12} f(p)$ , and note that this function is subexponential in  $k$  because its number of summands is linear in  $k$ . Because  $\lambda > \bar{c}$  and  $p \geq 2k + 2$ , we have that:

$$\pi(S_1) \leq \frac{f_1(k) \left(\frac{\bar{c}}{\lambda}\right)^{2k+2}}{Z}.$$

By Lemma 5.45, there is a constant  $C_1 = \bar{c}^2/(\lambda^2 C)$  such that:

$$\pi(S_1) \leq \frac{f_1(k) \left(\frac{\bar{c}}{\lambda}\right)^{2k+2}}{C \left[(\lambda\gamma)^{-2\sqrt{3}\tan\frac{\theta}{2}}\right]^k} = C_1 f_1(k) \left(\frac{\bar{c}(\lambda\gamma)^{\sqrt{3}\tan\frac{\theta}{2}}}{\lambda}\right)^{2k}.$$

For all  $\theta < 2 \tan^{-1}(\log_{\lambda\gamma}(\lambda/\bar{c})/\sqrt{3})$ , the term in parentheses above is less than one:

$$\xi := \frac{\bar{c}(\lambda\gamma)^{\sqrt{3}\tan\frac{\theta}{2}}}{\lambda} < \frac{\bar{c}(\lambda\gamma)^{\log_{\lambda\gamma}\left(\frac{\lambda}{2+\sqrt{2}}\right)}}{\lambda} = 1.$$




Figure 5.21: The path of length  $k$  (bold) from vertex  $v_1$  to the first land location in line  $b$  considered in the proof of Lemma 5.47; this path is used to calculate the gap height  $k$  in terms of the gap depth  $q$ . By also considering the reflection of this path from  $v_2$  (solid line), we can calculate the distance between the two objects to be  $q + 2\lceil w \rceil + 3$  (Lemma 5.48).

It follows that there exists a constant  $\xi < 1$  and a subexponential function  $f_2(k) = C_1 f_1(k)$  such that  $\pi(S_1) < f_2(k)\xi^k$ , proving the theorem.  $\square$

Since  $n = 4k + 5$  by Lemma 5.42, the probability that the inner perimeter is above point  $m$  is also exponentially small in  $n$ , the number of particles.

As an example, for  $\lambda = 4$  and  $\gamma = 2$  (the parameters of the simulation in Figure 5.15), our methods give  $\theta_1 = 0.0879 \sim 5.03^\circ$ . However, simulations suggest this bound is far from tight. In general, as  $\lambda$  increases, so does the angle  $\theta_1$ : a stronger bias towards a shorter perimeter means the bridge forms closer to the bottom of the gap and at even larger angles the bridge remains below  $m$ . Similarly, as  $\gamma$  decreases the bridge moves down towards the bottom of the gap and at even larger angles remains below  $m$ .

### 5.8.2 Proofs for Large $\theta$

We now consider the set  $S_2 = \Omega \setminus S_1$ , which consists of all configurations where the inner perimeter goes through or below  $m$ . We will show that for some large angles  $\theta$ , for all  $\lambda > 2 + \sqrt{2}$  and  $\gamma > (2 + \sqrt{2})^4 \lambda^4$ ,  $\pi(S_2)$  is exponentially small. While a lower bound on  $\gamma$  is necessary for the proofs presented below, we believe this is an artifact of our proof rather

than the problem itself and suspect this requirement can be loosened or removed altogether.

For  $\theta \geq \pi/3$ , it is no longer true that a V-shaped land mass of height  $k$  has exactly  $k - 1$  lattice lines between  $b$  and  $e$ . We define a new quantity  $q$ , the *gap depth*, as the length of a shortest path from  $e$  to  $b$  in  $\Gamma$ ; unlike in the definition of the height  $k$  of a gap, this shortest path is not required to stay on land locations. The Euclidean distance between  $e$  and  $b$  is then  $\sqrt{3}q/2$ . Furthermore,  $q$  can be expressed as a function of  $k$  and  $\theta$ .

**Lemma 5.47.** *For a V-shaped land mass of height  $k$  and angle  $\theta \geq \pi/3$ , the gap depth  $q$  satisfies*

$$k = \left\lceil \left( \frac{1}{2} + \frac{\sqrt{3}}{2} \tan \frac{\theta}{2} \right) q \right\rceil.$$

*Proof.* Consider the path from  $v_1$  to line  $b$  that leaves  $v_1$  forming an angle of  $2\pi/3$  with  $e$ , and then proceeds along  $b$  until it reaches a land location; see Figure 5.21, where this path is shown in bold. The total length of this path is  $k$ , and its first segment from  $v_1$  to  $b$  is length  $q$ . Let  $w$  be the length of  $b$  between this path's turning point and  $\ell_1$ ; then  $k = q + \lceil w \rceil$ . This path and  $\ell_1$  form an obtuse triangle where two sides have lengths  $q$  and  $w$ , respectively. The angle opposite the side of length  $w$  is  $\theta/2 - \pi/6$ , while the angle opposite the side of length  $q$  is  $\pi - 2\pi/3 - (\theta/2 - \pi/6) = \pi/2 - \theta/2$ . Length  $w$  can be calculated in terms of length  $q$  with the law of sines:

$$w = \frac{\sin\left(\frac{\theta}{2} - \frac{\pi}{6}\right)}{\sin\left(\frac{\pi}{2} - \frac{\theta}{2}\right)} q = \frac{\sin\frac{\theta}{2} \cos\frac{\pi}{6} - \cos\frac{\theta}{2} \sin\frac{\pi}{6}}{\cos\frac{\theta}{2}} q = \frac{\frac{\sqrt{3}}{2} \sin\frac{\theta}{2} - \frac{1}{2} \cos\frac{\theta}{2}}{\cos\frac{\theta}{2}} q = \frac{q\sqrt{3}}{2} \tan\frac{\theta}{2} - \frac{q}{2}.$$

Because  $q$  is an integer, it follows that

$$k = q + \lceil w \rceil = \left\lceil q + \frac{q\sqrt{3}}{2} \tan\frac{\theta}{2} - \frac{q}{2} \right\rceil = \left\lceil \left( \frac{1}{2} + \frac{\sqrt{3}}{2} \tan\frac{\theta}{2} \right) q \right\rceil,$$

which is the desired result. □

For simplicity, we do the bulk of our analysis using  $q$  instead of  $k$ . The previous lemma shows that proving an expression is exponentially small in  $q$  implies it is also exponentially




Figure 5.22: From the proof of Lemma 5.49: (a) An example of a shortest path between land locations on opposite sides of the gap passing through midpoint  $m$ . (b) The four possible locations for midpoint  $m$  for which a shortest path passing through or below  $m$  contains  $m'$ , and a shortest path from  $m'$  to a land location (solid line).

small in  $k$ .

**Lemma 5.48.** *For any V-shaped land mass of gap depth  $q$  and angle  $\theta \geq \pi/3$ , any configuration  $\sigma$  has perimeter at least*

$$p(\sigma) \geq \left(2\sqrt{3} \tan \frac{\theta}{2}\right) q + 6.$$

*Proof.* We first bound the distance between the two objects on either side of the gap. Using the length  $w$  from the proof of Lemma 5.47, the distance between the two objects in any configuration is  $q + 2\lceil w \rceil + 3 \geq q + 2w + 3$  (see Figure 5.21). The perimeter of any particle configuration is at least twice this distance, so for any  $\sigma$ ,

$$p(\sigma) \geq 2q + 4w + 6 = 2q + 4 \left( \frac{q\sqrt{3}}{2} \tan \frac{\theta}{2} - \frac{q}{2} \right) + 6 = \left(2\sqrt{3} \tan \frac{\theta}{2}\right) q + 6,$$

which is the desired bound. □

**Lemma 5.49.** *For any V-shaped land mass of gap depth  $q$  and angle  $\theta > \pi/3$ , any configuration  $\sigma \in S_2$  (passing below or through midpoint  $m$  of the gap) has gap perimeter  $g(\sigma) \geq \frac{q}{2}$ .*

*Proof.* If  $\sigma \in S_2$ , i.e., if its inner perimeter passes through or below  $m$ , then it must contain



a path that starts and ends at land locations and also passes through or below  $m$ . We consider all such paths and give a lower bound on the number of gap locations they must contain. The shortest such paths start and end on opposite sides of the gap, so we focus on paths of this type.

If  $m$  is a vertex of  $\Gamma$ , one shortest path between land locations passing through  $m$  leaves  $m$  along the two lattice lines not parallel to  $e$  and follows them until reaching the land mass, as in Figure 5.22a. If  $m$  is on a lattice edge, a shortest path passing below  $m$  is constructed in the same way, beginning from each of the edge's endpoints. Otherwise, if  $m$  is neither a lattice point nor on a lattice edge, the same procedure is followed for the first lattice point or lattice edge below  $m$ . In all cases, let  $m'$  be the point of intersection between this path and  $\ell^*$ , the line perpendicular to  $e$  through  $v_1$ . Figure 5.22b shows all the possible locations of  $m$  producing a particular  $m'$ . Inspection shows that in all of these cases,  $m'$  is contained in the  $2\lfloor \frac{q+1}{4} \rfloor$ -th lattice line below  $e$ .

Let  $\bar{\ell}_1$  be the line from  $v_1$  to  $b$  forming an angle of  $2\pi/3$  with  $e$ ; see Figure 5.22b. Because  $\theta > \pi/3$ , all vertices of  $\Gamma$  contained in  $\bar{\ell}_1$ , except  $v_1$ , are gap locations. Any shortest path from  $m'$  to a land location must share a vertex of  $\Gamma$  with line  $\bar{\ell}_1$ . Because  $m'$  is in the  $2\lfloor \frac{q+1}{4} \rfloor$ -th lattice line below  $e$ , any path from  $m'$  to  $\bar{\ell}_1$  is of length at least  $\lfloor \frac{q+1}{4} \rfloor$  and contains at least  $\lfloor \frac{q+1}{4} \rfloor + 1$  gap locations, including both of its endpoints. By symmetry, this means any path between land locations passing below  $m$ , and thus any inner perimeter of a particle configuration passing below  $m$ , contains at least

$$2 \left( \left\lfloor \frac{q+1}{4} \right\rfloor + 1 \right) \geq 2 \left( \frac{q-2}{4} + 1 \right) \geq \frac{q}{2}$$

gap locations, as claimed. □

**Theorem 5.50.** *Let  $\lambda > 2 + \sqrt{2} =: \bar{c}$  and  $\gamma > (\lambda \bar{c})^4$ . Then there exists a constant  $\theta_2 > \pi/3$  such that for all V-shaped land masses with angle  $\theta \in (\pi/3, \theta_2)$ , the probability that the inner perimeter goes through or below midpoint  $m$  is exponentially small in  $k$ , the height*

of the gap.

*Proof.* Recall  $S_2$  is the set of all configurations whose inner perimeter goes through or below  $m$ . We show that  $\pi(S_2)$  is exponentially small in  $k$ , the height of the gap. By definition,

$$\pi(S_2) = \frac{\sum_{\sigma \in S_2} \lambda^{-p(\sigma)} \gamma^{-g(\sigma)}}{Z}.$$

By Lemma 5.42, the number of particles and objects in  $\sigma_0$  for a land mass of height  $k$  is  $4k + 7$ . Since  $\sigma_0$  is a path of width 2 and every particle occupies a land location,  $p(\sigma_0) = 4k + 7$  and  $g(\sigma_0) = 0$ . Thus,

$$Z = \sum_{\sigma \in \Omega} \lambda^{-p(\sigma)} \gamma^{-g(\sigma)} \geq \lambda^{-p(\sigma_0)} \gamma^{-g(\sigma_0)} = \lambda^{-4k-7}.$$

It is simpler to work with gap depth  $q$  instead of gap height  $k$ . By Lemma 5.47,  $k$  satisfies  $k \leq \left(\frac{1}{2} + \frac{\sqrt{3}}{2} \tan \frac{\theta}{2}\right) q + 1$ , so

$$Z \geq \lambda^{-4k-7} \geq \lambda^{-4\left(\frac{1}{2} + \frac{\sqrt{3}}{2} \tan \frac{\theta}{2}\right)q - 4 - 7} = \lambda^{-(2+2\sqrt{3} \tan \frac{\theta}{2})q - 11}.$$

Combining this with Lemma 5.49,

$$\pi(S_2) = \sum_{\sigma \in S_2} \frac{\lambda^{-p(\sigma)} \gamma^{-g(\sigma)}}{Z} \leq \lambda^{(2+2\sqrt{3} \tan \frac{\theta}{2})q + 11} \sum_{\sigma \in S_2} \lambda^{-p(\sigma)} \gamma^{-\frac{q}{2}}.$$

Let  $p_{min}$  (resp.,  $p_{max}$ ) be the minimum (resp., maximum) possible perimeter for a valid particle configuration in  $S_2$ . By Lemma 5.48,  $p_{min} \geq 2\sqrt{3} \tan(\theta/2)q$ . As shown in the proof of Theorem 5.46,  $p_{max} = 8k + 12$ ; in terms of  $q$ , by Lemma 5.47,

$$p_{max} \leq 8 \left( \frac{q}{2} + \frac{q\sqrt{3}}{2} \tan \frac{\theta}{2} + 1 \right) + 12 = 4q + 4q\sqrt{3} \tan \frac{\theta}{2} + 20.$$

Using the result from [21] which upper bounds the number of particle configurations with perimeter  $p$  by the expression  $f(p)\bar{c}^p$ , for some subexponential function  $f$ , we have

that:

$$\begin{aligned}
\pi(S_2) &\leq \lambda^{(2+2\sqrt{3}\tan\frac{\theta}{2})q+11} \sum_{p=p_{min}}^{p_{max}} f(p) \bar{c}^p \lambda^{-p} \gamma^{-\frac{q}{2}} \\
&\leq \lambda^{(2+2\sqrt{3}\tan\frac{\theta}{2})q+11} \left( \sum_{p=p_{min}}^{p_{max}} f(p) \right) \left( \frac{\bar{c}}{\lambda} \right)^{p_{min}} \gamma^{-\frac{q}{2}} \\
&\leq \left( \lambda^{11} \sum_{p=p_{min}}^{p_{max}} f(p) \right) \left( \lambda^{(2+2\sqrt{3}\tan\frac{\theta}{2})} \left( \frac{\bar{c}}{\lambda} \right)^{2\sqrt{3}\tan\frac{\theta}{2}} \gamma^{-\frac{1}{2}} \right)^q \\
&= \left( \lambda^{11} \sum_{p=p_{min}}^{p_{max}} f(p) \right) \left( \lambda^2 \bar{c}^{2\sqrt{3}\tan\frac{\theta}{2}} \gamma^{-\frac{1}{2}} \right)^q.
\end{aligned}$$

The first parentheses is a function  $f_1(q)$  that is subexponential in  $q$ , as it has a polynomial number of summands based on our calculations of  $p_{min}$  and  $p_{max}$  (which are expressions in terms of  $q$ ), and each summand is subexponential. When the term in the second set of parentheses above is less than one, this whole expression is exponentially small in  $q$ , the gap depth, and thus exponentially small in the gap height  $k$ . This holds whenever  $\theta$  satisfies:

$$\theta < 2 \tan^{-1} \left( \frac{1}{2\sqrt{3}} \log_{\bar{c}} (\gamma^{1/2} \lambda^{-2}) \right) = 2 \tan^{-1} \left( \frac{1}{\sqrt{3}} \log_{\bar{c}} \left( \frac{\gamma^{1/4}}{\lambda} \right) \right) =: \theta_2.$$

Whenever  $\gamma^{1/4}/\lambda > \bar{c}$  — i.e., whenever  $\gamma > (\lambda\bar{c})^4$  — the argument of  $\tan^{-1}$  above is at least  $1/\sqrt{3}$ , and thus  $\theta_2 > \pi/3$ . It follows that whenever  $\gamma > (\lambda\bar{c})^4$  and  $\theta \in (\pi/3, \theta_2)$ ,

$$\pi(S_2) < f_1(q) \psi^P,$$

where  $f_1(q)$  is a function that is subexponential in  $q$  and  $\psi < 1$ . We conclude that under these conditions the weight of set  $S_2$  at stationarity is exponentially small in  $q$ , and thus also exponentially small in  $k$ , the gap height.  $\square$

This concludes our work on the shortcut bridging problem, which we have provably solved with a Markov chain  $\mathcal{M}$  that can be directly translated to a stochastic, distributed,

local, asynchronous algorithm  $\mathcal{A}$ . Furthermore, in the special case of bridging over the gap in a V-shaped land mass, we rigorously analyzed the effect of the gap's internal angle, showing that below one threshold angle the bridge will shortcut near the bottom of the gap, and above another threshold angle the bridge will remain close to land, with high probability.

The successful application of our stochastic approach to shortcut bridging suggests that this approach may be useful for other types of bridging problems as well; one related behavior of particular interest is “exploration bridging”, where a particle system first explores its environment to discover sites of interest, and then converges to a bridge-like structure between these sites. We are also interested in formulating alternative local rules for shortcut bridging which result in bridges that appear more “structurally sound”, though we suspect that the information needed for doing so may be difficult to encode in our particle systems due to the constant-size memory constraint of the amoebot model.

## **5.9 Further Applications of the Stochastic Approach to Programmable Matter**

Beyond bridging, compression, and expansion, there exists a plethora of other problems within programmable matter for which the stochastic approach seems promising; we believe we have so far only scratched the surface. For example, in more recent applied work in swarm robotics, we use a Markov chain algorithm for a self-organizing particle system to provide a theoretical explanation of a behavior our physicist collaborators were observing in their robot swarm [23]. Current work includes applying the stochastic approach to the *separation* problem, where particles of different colors either intermingle or segregate based on the values of bias parameters.

More broadly, for any objective that can be described by a global energy function, where changes to the energy function when one particle moves one unit can be calculated with local information, our approach offers a method for developing a distributed algorithm that converges to a distribution that favors configurations with the lowest energy. However,

entropy concerns must also be taken into account. If there are many more configurations with high energy than with low energy, even though the low energy configurations are most favorable a configuration drawn from the stationary distribution may be more likely to have high energy. This is why, for instance, we don't see compression for all  $\lambda > 1$ , even though  $\lambda > 1$  corresponds to favoring smaller perimeter configurations. For this reason, chains must be biased enough to guarantee low energy configurations dominate the state space, even if there are many high energy configurations. We used Peierls arguments to analyze this energy/entropy trade-off for compression, expansion, and the dependence of bridge structure on gap angle, and expect similar approaches may work for future problems.

## REFERENCES

- [1] L. Adleman. “Molecular computation of solutions to combinatorial problems”. In: *Science* 266 (5187) (1994), pp. 1021–1024.
- [2] D. Aldous and P. Diaconis. “Shuffling Cards and Stopping Times”. In: *The American Mathematical Monthly* 93 (5) (1986), pp. 333–348.
- [3] American Association for Public Opinion Research. *Sampling Methods for Political Polling*. Accessed 2/27/2018. URL: <http://www.aapor.org/Education-Resources/Election-Polling-Resources/Sampling-Methods-for-Political-Polling.aspx>.
- [4] M. Andrés Arroyo, S. Cannon, J. Daymude, D. Randall, and A. Richa. “A Stochastic Approach to Shortcut Bridging in Programmable Matter”. In: *23rd International Conference on DNA Computing and Molecular Programming (DNA23)*. 2017, pp. 122–138.
- [5] E. Bampas, J. Czyzowicz, L. Gąsieniec, D. Ilcinkas, and A. Labourel. “Almost Optimal Asynchronous Rendezvous in Infinite Multidimensional Grids”. In: *Distributed Computing: 24th International Symposium, DISC 2010*. 2010, pp. 297–311.
- [6] R. Bauerschmidt, H. Duminil-Copin, J. Goodman, and G. Slade. “Lectures on Self-Avoiding Walks”. In: *Clay Mathematics Proceedings* 15 (2012).
- [7] R. J. Baxter, I. G. Enting, and S. K. Tsang. “Hard-square lattice gas”. In: *Journal of Statistical Physics* 22 (4) (1980), pp. 465–489.
- [8] D. Bayer and P. Diaconis. “Trailing the Dovetail Shuffle to its Lair”. In: *The Annals of Applied Probability* 2 (2) (May 1992), pp. 294–313.
- [9] I. Benjamini, N. Berger, C. Hoffman, and E. Mossel. “Mixing Times of the Biased Card Shuffling and the Asymmetric Exclusion Process”. In: *Transactions of the American Mathematical Society* 357 (8) (2005), pp. 3013–3029.
- [10] P. Bhakta, S. Miracle, D. Randall, and A. P. Streib. “Mixing Times of Markov Chains for Self-Organizing Lists and Biased Permutations”. In: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. 2013, pp. 1–15.

- [11] N. Bhatnagar, S. Greenberg, and D. Randall. “The effects of boundary conditions on mixing rates of Markov chains”. In: *Approximation, Randomization and Combinatorial Optimization Algorithms and Techniques* (2006), pp. 280–291.
- [12] A. Blanca, D. Galvin, D. Randall, and P. Tetali. “Phase Coexistence and Slow Mixing for the Hard-Core Model on  $\mathbb{Z}^2$ ”. In: *17th International Workshop on Randomization and Computation (RANDOM ’13)*. 2013, pp. 379–394.
- [13] A. Blanca, P. Caputo, A. Sinclair, and E. Vigoda. “Spatial Mixing and Non-local Markov chains”. In: *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2018, pp. 1965–1980.
- [14] C. Borgs, J. Chayes, A. Frieze, J. Kim, P. Tetali, E. Vigoda, and V. Vu. “Tortoise Mixing of Some MCMC Algorithms in Statistical Physics”. In: *40th IEEE Symp. on Found. of Comp. Sci. FOCS 1999*. 1999, pp. 218–229.
- [15] E. Boros and Z. Füredi. “Rectangular Dissections of a Square”. In: *European Journal of Combinatorics* 9 (1988), pp. 271–280.
- [16] R. Bubley and M. Dyer. “Path coupling: A technique for proving rapid mixing in Markov chains”. In: *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS)*. 1997, pp. 223–231.
- [17] R. Bubley, M. Dyer, C. Greenhill, and M. Jerrum. “Approximately counting colorings of small degree graphs”. In: *SIAM Journal of Computing* 29 (1999), pp. 387–400.
- [18] M. Burr, S. W. Choi, B. Galehouse, and C. K. Yap. “Complete subdivision algorithms, II: Isotopic meshing of singular algebraic curves”. In: *Journal of Symbolic Computation* 47 (2) (2012), pp. 131–152.
- [19] S. Camazine, K. Visscher, J. Finley, and S. Vetter. “House-Hunting by Honey Bee Swarms: Collective Decisions and Individual Behaviors”. In: *Insectes Sociaux* 46 (), pp. 348–360.
- [20] S. Cannon, D. Levin, and A. Stauffer. “Polynomial mixing of the edge-flip Markov chain for Unbiased Dyadic Tilings”. In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017)*. 2017, 34:1–34:21.
- [21] S. Cannon, J. J. Daymude, D. Randall, and A. W. Richa. “A Markov Chain Algorithm for Compression in Self-Organizing Particle Systems”. In: *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC ’16)*. 2016, pp. 279–288.

- [22] S. Cannon, J. Daymude, C. Gokmen, D. Randall, and A. Richa. “Brief Announcement: A Local Stochastic Algorithm for Separation in Heterogeneous Self-Organizing Particle Systems”. Under Review.
- [23] S. Cannon, J. Daymude, D. Goldman, S. Li, D. Randall, A. Richa, W. Savoie, and R. Warkentin. “Phototactic Supersmarticles”. In: *The 2nd International Symposium on Swarm Behavior and Bio-Inspired Robotics (SWARM)*. 2017, pp. 377–384.
- [24] S. Cannon, S. Miracle, and D. Randall. “Phase Transitions in random dyadic tilings and rectangular dissections”. In: *Proceedings of the 26th Symposium on Discrete Algorithms (SODA)*. 2015.
- [25] S. Cannon and D. Randall. “Sampling on Lattices with Free Boundary Conditions Using Randomized Extensions”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by R. Krauthgamer. Society for Industrial and Applied Mathematics, 2016.
- [26] P. Caputo, F. Martinelli, A. Sinclair, and A. Stauffer. “Dynamics of lattice triangulations on thin rectangles”. In: *Electronic Journal of Probability* 21 (29) (2016).
- [27] P. Caputo, F. Martinelli, A. Sinclair, and A. Stauffer. “Random Lattice Triangulations: Structure and Algorithms”. In: *The Annals of Applied Probability* 25 (4) (2015), pp. 1650–1685.
- [28] F. Cesi. “Quasi-factorization of the entropy and logarithmic Sobolev inequalities for Gibbs random fields”. In: *Probability Theory and Related Fields* 120 (4) (2001), pp. 569–584.
- [29] M.-F. Chen. “Trilogy of couplings and general formulas for lower bound of spectral gap”. In: *Probability towards 2000 (New York, 1995)*. Vol. 128. Lecture Notes in Statistics. Springer, New York, 1998, pp. 123–136.
- [30] K. C. Cheung, E. D. Demaine, J. R. Bachrach, and S. Griffith. “Programmable Assembly With Universally Foldable Strings (Moteins)”. In: *IEEE Transactions on Robotics* 27 (4) (2011), pp. 718–729.
- [31] H. Cohn, M. Larsen, and J. Propp. “The Shape of a Typical Boxed Plane Partition”. In: *New York Journal of Mathematics* 4 (1998), pp. 137–165.
- [32] B. Cousins and S. Vempala. “Bypassing KLS: Gaussian Cooling and an  $\mathcal{O}^*(n^3)$  Volume Algorithm”. In: *STOC ’15: Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*. 2015.



- [33] P. Cuff, J. Ding, O. Louidor, E. Lubetzky, Y. Peres, and A. Sly. “Glauber Dynamics for the Mean-Field Potts Model.” In: *Journal of Statistical Physics* 149 (3) (2012), pp. 432–477.
- [34] C. David, P. Athina, G. Christophe, J. Nynika, R. Steffen, M. Achim, and T. Skylar. “3D-Printed Wood: Programming Hygroscopic Material Transformations”. In: *3D Printing and Additive Manufacturing* 2 (3) (2015), pp. 106–116.
- [35] J. Daymude, R. Gmyr, A. Richa, C. Scheideler, and T. Strothmann. “Improved Leader Election for Self-Organizing Programmable Matter”. In: *Algorithms for Sensor Systems (ALGOSENSORS ’17)*. 2017, pp. 127–140.
- [36] J. J. Daymude, Z. Derakhshandeh, R. Gmyr, A. Porter, A. W. Richa, C. Scheideler, and T. Strothmann. “On the Runtime of Universal Coating for Programmable Matter”. In: *Natural Computing* 17 (1) (2017), pp. 81–96.
- [37] G. DeMicheli, P. Antognetti, and A. Sangiovanni-Vincentelli, eds. *Design systems for VLSI circuits: logic synthesis and silicon compilation*. Springer, 1987.
- [38] Z. Derakhshandeh, R. Gmyr, A. Richa, C. Scheideler, and T. Strothmann. “An Algorithmic Framework for Shape Formation Problems in Self-Organizing Particle Systems”. In: *Proc. of the 2nd Ann. Int. Conf. on Nanoscale Computing and Comm., NANOCOM’ 15*. 2015, 21:1–21:2.
- [39] Z. Derakhshandeh, S. Dolev, R. Gmyr, A. Richa, C. Scheideler, and T. Strothmann. “Brief Announcement: Amoebot - a New Model for Programmable Matter”. In: *26th ACM Symp. on Parallelism in Alg. and Arch., SPAA ’14*. 2014, pp. 220–222.
- [40] Z. Derakhshandeh, R. Gmyr, T. Strothmann, R. A. Bazzi, A. W. Richa, and C. Scheideler. “Leader Election and Shape Formation with Self-organizing Programmable Matter”. In: *Proceedings of the 21st International Conference on DNA Computing and Molecular Programming (DNA21)*. 2015, pp. 117–132.
- [41] Z. Derakhshandeh, R. Gmyr, A. Richa, C. Scheideler, and T. Strothmann. “Universal Coating for Programmable Matter”. In: *Theoretical Computer Science* 671 (2017), pp. 56–68.
- [42] Z. Derakhshandeh, R. Gmyr, A. W. Richa, C. Scheideler, and T. Strothmann. “Universal Shape Formation for Programmable Matter”. In: *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*. SPAA ’16. 2016, pp. 289–299.
- [43] P. Diaconis and L. Saloff-Coste. “Comparison theorems for reversible Markov chains”. In: *The Annals of Applied Probability* 3 (1993), pp. 696–730.

- [44] P. Diaconis and D. Stroock. “Geometric Bounds for Eigenvalues of Markov Chains”. In: *The Annals of Applied Probability* 1 (1) (1991), pp. 36–61.
- [45] P. Diaconis. *Group Representations in Probability and Statistics*. Vol. 11. LECTURE NOTES-MONOGRAPH SERIES. Institute of Mathematical Statistics, 1988.
- [46] P. Diaconis and M. Shahshahani. “Generating a random permutation with random transpositions”. In: *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 57 (2) (1981), pp. 159–179.
- [47] J. Ding, E. Lubetzky, and Y. Peres. “Mixing Time of Critical Ising Model on Trees is Polynomial in the Height”. In: *Communications in Mathematical Physics* 295 (1) (2010), pp. 161–207.
- [48] J. Ding, E. Lubetzky, and Y. Peres. “The Mixing Time Evolution of Glauber Dynamics for the Mean-Field Ising Model”. In: *Communications in Mathematical Physics* 289 (2) (2009), pp. 725–764.
- [49] R. Dobrushin. “The problem of Uniqueness of a Gibbsian Random Field and the Problem of Phase Transitions”. In: *Functional Analysis and Its Applications* 2 (1968), pp. 302–312.
- [50] S. Douglas, H. Dietz, T. Liedl, B. Högberg, F. Graf, and W. Shih. “Self-assembly of DNA into nanoscale three-dimensional shapes”. In: *Nature* 459 (2009), pp. 414–418.
- [51] H. Duminil-Copin and S. Smirnov. “The connective constant of the honeycomb lattice equals  $\sqrt{2 + \sqrt{2}}$ ”. In: *Annals of Mathematics* 175 (3) (2012), pp. 1653–1665.
- [52] D. Durfee, J. Peebles, R. Peng, and A. B. Rao. “Determinant-Preserving Sparsification of SDDM Matrices with Applications to Counting and Sampling Spanning Trees”. In: *58th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2017, pp. 926–937.
- [53] M. Dyer and C. Greenhill. “A more rapidly mixing Markov chain for graph colorings”. In: *Random Structures & Algorithms* 13 (3-4) (1998), pp. 285–317.
- [54] M. Dyer, A. Sinclair, E. Vigoda, and D. Weitz. “Mixing in time and space for lattice spin systems: A combinatorial view”. In: *Random Structures & Algorithms* 24 (4) (2004), pp. 461–479.
- [55] W. Feller. *An Introduction to Probability Theory and Its Applications*. Vol. 1. Wiley, 1968.

- [56] P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer. “Arbitrary Pattern Formation by Asynchronous, Anonymous, Oblivious Robots”. In: *Theoretical Computer Science* 407 (2008), pp. 412–447.
- [57] P. Flocchini, G. Prencipe, and N. Santoro. “Distributed Computing by Oblivious Mobile Robots”. In: *Synthesis Lectures on Distributed Computing Theory* 3 (2) (2012), pp. 1–185.
- [58] S. Fortune. “A Sweepline Algorithm for Voronoi Diagrams”. In: *Proceedings of the Second Annual Symposium on Computational Geometry*. SCG ’86. Yorktown Heights, New York, USA: ACM, 1986, pp. 313–322.
- [59] A. Frieze and E. Vigoda. “A survey on the use of Markov chains to randomly sample colourings”. In: *Combinatorics, Complexity and Chance* (2007).
- [60] R. Gheissari and E. Lubetzky. “Mixing times of critical 2D Potts models”. Submitted. Available at <https://arxiv.org/abs/1607.02182>.
- [61] L. A. Goldberg, R. Martin, and M. Paterson. “Random Sampling of 3-Colorings in  $\mathbb{Z}^2$ ”. In: *Random Structures and Algorithms*. Vol. 24. 2004, pp. 279–302.
- [62] S. Greenberg, A. Pascoe, and D. Randall. “Sampling Biased Lattice Configurations using Exponential Metrics”. In: *20th Symposium on Discrete Algorithms*. Vol. 21. 2009, pp. 225–251.
- [63] L. Grilz, Z. Gao, and W. Schmitt. “Statistical Analysis of Chemical Transformation Kinetics Using Markov-Chain Monte Carlo Methods”. In: *Environmental Science & Technology* 45 (10) (2011), pp. 4429–4437.
- [64] R. Haggkvist, P. Lindberg, and B. Lindstrom. “Dissecting a square into rectangles of equal area”. In: *Discrete Mathematics* 47 (1983), pp. 321–323.
- [65] T. Hayes and E. Vigoda. “A non-Markovian coupling for randomly sampling colorings”. In: *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*. 2003, pp. 618–627.
- [66] E. Ising. “Contribution to the Theory of Ferromagnetism”. In: *Z. Phys.* 31 (1925), pp. 253–258.
- [67] S. Janson, D. Randall, and J. Spencer. “Random Dyadic Tilings of the Unit Square”. In: *Random Structures and Algorithms* 21 (3-4 2002), pp. 225–251.
- [68] R. Jeanson, C. Rivault, J. Deneubourg, S. Blanco, R. Fournier, C. Jost, and G. Theraulaz. “Self-Organized Aggregation in Cockroaches”. In: *Animal Behaviour* 69 (2005), pp. 169–180.

- [69] I. Jensen. “A parallel algorithm for the enumeration of benzenoid hydrocarbons”. In: *Journal of Statistical Mechanics: Theory and Experiment* 2009 (02) (2009), P02065.
- [70] I. Jensen. “Improved lower bounds on the connective constants for two-dimensional self-avoiding walks”. In: *Journal of Physics A: Mathematical and General* 37 (48) (2004), p. 11521.
- [71] M. Jerrum, L. Valiant, and V. Vazirani. “Random generation of combinatorial structures from a uniform distribution”. In: *Theoretical Computer Science* 43 (1986), pp. 169–188.
- [72] M. Johnson and T. L. Griffiths. “Bayesian inference for PCFGs via Markov chain Monte Carlo”. In: *In Proceedings of the North American Conference on Computational Linguistics (NAACL 07. 2007.*
- [73] J. C. Lagarias, J. H. Spencer, and J. P. Vinson. “Counting Dyadic Equipartitions of the Unit Square”. In: *Discrete Mathematics* 257 (2-3) (Nov. 2002), pp. 481–499.
- [74] B Larget and D. Simon. “Markov Chasin Monte Carlo Algorithms for the Bayesian Analysis of Phylogenetic Trees”. In: *Molecular Biology and Evolution* 16 (6) (1999), p. 750.
- [75] L. Lau and M. Molloy. “Randomly Colouring Graphs with Girth Five and Large Maximum Degree”. In: *LATIN 2006: Theoretical Informatics*. Ed. by J. Correa, A. Hevia, and M. Kiwi. Vol. 3887. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2006, pp. 665–676.
- [76] D. Levin, Y. Peres, and E. Wilmer. *Markov chains and mixing times*. With a chapter on coupling from the past by J.G. Propp and D.B. Wilson. Providence, R.I. American Mathematical Society, 2009.
- [77] D. A. Levin, M. J. Luczak, and Y. Peres. “Glauber dynamics for the mean-field Ising model: cut-off, critical power law, and metastability”. In: *Probab. Theory Related Fields* 146 (1-2) (2010), pp. 223–265.
- [78] E. Lubetzky and A. Sly. “Critical Ising on the Square Lattice Mixes in Polynomial Time”. In: *Communications in Mathematical Physics* 313 (3) (2012), pp. 815–836.
- [79] M. Luby, D. Randall, and A. Sinclair. “Markov chain algorithms for planar lattice structures”. In: *SIAM Journal on Computing* 31 (2001), pp. 167–192.
- [80] N. Lynch. *Distributed Algorithms*. Morgan Kauffman, 1996.

- [81] N. Madras and D. Randall. “Markov chain decomposition for convergence rate analysis”. In: *The Annals of Applied Probability* 12 (2) (2002), pp. 581–606.
- [82] R. Martin and D. Randall. “Pfaffian algorithms for sampling routings on regions with free boundary conditions”. In: *3rd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM) in Lecture Notes in Computer Science* 1671 (1999), pp. 257–268.
- [83] R. Martin and D. Randall. “Sampling Adsorbing Staircase Walks Using a New Markov Chain Decomposition Method”. In: *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS)*. 2000, pp. 492–502.
- [84] F. Martinelli, A. Sinclair, and D. Weitz. “Fast mixing for independent sets, colorings and other models on trees”. In: *Proc. 15th ACM/SIAM Symposium on Discrete Algorithms* (2004), pp. 449–458.
- [85] F. Martinelli, A. Sinclair, and D. Weitz. “The Ising model on trees: Boundary conditions and mixing time”. In: *Comm. Mathematical Physics* 250 (2004), 301334.
- [86] F. Martinelli. “Lectures on Glauber Dynamics for Discrete Spin Models”. In: *Lectures on Probability Theory and Statistics: Ecole d’Eté de Probabilités de Saint-Flour XXVII - 1997*. Ed. by P. Bernard. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 93–191.
- [87] L. McShine and P. Tetali. “On the mixing time of the triangulation walk and other Catalan structures”. In: *DIMACS-AMS Volume on Randomization Methods in Algorithm Design* 43 (1998), pp. 147–160.
- [88] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. “Equation of State Calculations by Fast Computing Machines”. In: *Journal of Chemical Physics* 21 (1953), pp. 1087–1092.
- [89] N. Mlot, C. Tovey, and D. Hu. “Fire Ants Self- Assemble into Waterproof Rafts to Survive Floods”. In: *Proc. of the National Academy of Sci.* 108 (2011), pp. 7669–7673.
- [90] A. Mohammed, P. Šulc, J. Zenk, and R. Schulman. “Self-assembling DNA nanotubes to connect molecular landmarks”. In: *Nature Nanotechnology* 12 (2017), pp. 312–316.
- [91] M. Molloy. “The Glauber Dynamics on Colourings of a Graph with High Girth and Maximum Degree”. In: *Proceedings of the Thirty-fourth Annual ACM Symposium on Theory of Computing*. STOC ’02. Montreal, Quebec, Canada: ACM, 2002, pp. 91–98.

- [92] M. Molloy, B. Reed, and W. Steiger. “On the mixing rate of the triangulation walk”. In: *DIMACS-AMS Volume on Randomization Methods in Algorithm Design* 43 (1998), pp. 179–190.
- [93] P. Moubarak and P. Ben-Tzvi. “Modular and reconfigurable mobile robotics”. In: *Robotics and Autonomous Systems* 60 (12) (2012), pp. 1648–1663.
- [94] H. Murata, K. Fijiyoshi, T. Watanabe, and Y. Kajitani. “A mapping from sequence-pair to rectangular dissection”. In: *Design Automation Conference*. 1997, pp. 625–633.
- [95] T. Nakagaki, H. Yamada, and Ágota Tóth. “Intelligence: Maze-solving by an amoeboid organism”. In: *Nature* 407 (2000), p. 470.
- [96] A. Pönitz and P. Tittmann. “Improved Upper Bounds for Self-Avoiding Walks in  $\mathbb{Z}^d$ ”. In: *The Electronic Journal of Combinatorics* 7 (R21) (2000).
- [97] D. Randall and P. Tetali. “Analyzing Glauber dynamics by comparison of Markov chains”. In: *J. Mathematical Physics* 41 (2000), pp. 1598–1615.
- [98] C. R. Reid, M. J. Lutz, S. Powell, A. B. Kao, I. D. Couzin, and S. Garnier. “Army ants dynamically adjust living bridges in response to a cost–benefit trade-off”. In: *Proceedings of the National Academy of Sciences* 112 (49) (2015), pp. 15113–15118.
- [99] C. R. Reid, H. MacDonald, R. P. Mann, J. A. R. Marshall, T. Latty, and S. Garnier. “Decision-making without a brain: how an amoeboid organism solves the two-armed bandit”. In: *Journal of The Royal Society Interface* 13 (119) (2016).
- [100] R. Restrepo, J. Shin, P. Tetali, E. Vigoda, and L. Yang. “Improving Mixing Conditions on the Grid for Counting and Sampling Independent Sets”. In: *Probability Theory and Related Fields* 156 (2013), pp. 75–99.
- [101] C. Rivault and A. Cloarec. “Cockroach Aggregation: Discrimination between Strain Odours in *Blattella Germanica*”. In: *Animal Behaviour* 55 (1998), pp. 177–184.
- [102] M. Rubenstein, A. Cornejo, and R. Nagpal. “Programmable self-assembly in a thousand-robot swarm”. In: *Science* 345 (6198) (2014), pp. 795–799.
- [103] R. Salakhutdinov and G. Hinton. “Deep Boltzmann Machines”. In: *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2009.
- [104] C. Scott and R. D. Nowak. “Minimax-Optimal Classification With Dyadic Decision Trees”. In: *IEEE Transactions on Information Theory* 52 (4) (2006).

- [105] M. I. Shamos and D. Hoey. “Geometric Intersection Problems”. In: *Proceedings of the 17th Annual Symposium on Foundations of Computer Science*. SFCS '76. Washington, DC, USA: IEEE Computer Society, 1976, pp. 208–215.
- [106] A. Sinclair. *Algorithms for Random Generation & Counting: a Markov Chain Approach*. Birkhäuser, Boston, 1993.
- [107] A. Sinclair. “Improved bounds for mixing rates of Markov chains and multicommodity flow”. In: *Combinatorics, Probability and Computing* 1 (1992), pp. 351–370.
- [108] B. Sparrman, C. Matthews, S. Kernizan, A. Chadwick, N. Thomas, J. Laucks, and S. Tibbits. “Large-Scale Lightweight Transformable Structures”. In: *Proceedings of the 37th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*. 2017, pp. 572–581.
- [109] R. Thakker, A. Kamat, S. Bharambe, S. Chiddarwar, and K. M. Bhurchandi. “Re-BiS - Reconfigurable Bipedal Snake robot”. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2014, pp. 309–314.
- [110] L. Thomas. “Bounds on the Mass Gap for Finite Volume Stochastic Ising Models at Low Temperature”. In: *Comm. Math. Phys.* 126 (1989), pp. 1–11.
- [111] L. Thomas. “Quantum Heisenberg ferromagnets and stochastic exclusion processes”. In: *Journal of Mathematical Physics* 21 (7) (1980), pp. 1921–1924.
- [112] W. Tutte. “A Census of Planar Maps”. In: *Canad. J. Math.* 15 (1963), pp. 249–271.
- [113] J. C. Vera, E. Vigoda, and L. Yang. “Improved Bounds on the Phase Transition for the Hard-Core Model in 2 Dimensions”. In: *SIAM Journal on Discrete Mathematics* 29 (4) (2015), pp. 1895–1915.
- [114] E. Vigoda. “Improved bounds for sampling colorings”. In: *J. Mathematical Physics* 41 (3) (2000), pp. 1555–1569.
- [115] J. Wang and H. Zhou. “Linear Constraint Graph for Floorplan Optimization with Soft Blocks”. In: *International Conference on Computer-Aided Design (ICCAD'08)*. 2008, pp. 9–15.
- [116] B. Wei, M. Dai, and P. Yin. “Complex shapes self-assembled from single-stranded DNA tiles”. In: *Nature* 485 (2012), pp. 623–626.
- [117] E. Welzl. “The number of (and random) triangulations of planar point sets”. In: *Presentation at Oberwolfach Workshop on Combinatorics, Probability and Computing* (2006).

- [118] D. Wilson. “Mixing Times of Lozenge Tiling and Card Shuffling Markov Chains”. In: *Annals of Applied Probability* 14 (2004), pp. 274–325.
- [119] D. Woods, H. Chen, S. Goodfriend, N. Dabby, E. Winfree, and P. Yin. “Active Self-Assembly of Algorithmic Shapes and Patterns in Polylogarithmic Time”. In: *Proc. of the 4th Conf. on Innovations in Theoretical Computer Science*. 2013, pp. 353–354.
- [120] G. Zobrist. *Routing, Placement, and Partitioning*. Intellect Ltd., 1994.