A Local Stochastic Algorithm for Separation in Heterogeneous Self-Organizing Particle Systems

Sarah Cannon
Claremont McKenna College, Claremont, CA, USA
scannon@cmc.edu

Joshua J. Daymude
Computer Science, CIDSE, Arizona State University, Tempe, AZ, USA
jdaymude@asu.edu

Cem Gökmen
Georgia Institute of Technology, Atlanta, GA, USA
cgokmen@gatech.edu

Dana Randall
Georgia Institute of Technology, Atlanta, GA, USA
randall@cc.gatech.edu

Andréa W. Richa
Computer Science, CIDSE, Arizona State University, Tempe, AZ, USA
aricha@asu.edu

Abstract

We present and rigorously analyze the behavior of a distributed, stochastic algorithm for *separation* and *integration* in *self-organizing particle systems*, an abstraction of programmable matter. Such systems are composed of individual computational *particles* with limited memory, strictly local communication abilities, and modest computational power. We consider *heterogeneous* particle systems of two different colors and prove that these systems can collectively *separate* into different color classes or *integrate*, indifferent to color. We accomplish both behaviors with the same fully distributed, local, stochastic algorithm. Achieving separation or integration depends only on a single global parameter determining whether particles prefer to be next to other particles of the same color or not; this parameter is meant to represent external, environmental influences on the particle system. The algorithm is a generalization of a previous distributed, stochastic algorithm for *compression* (PODC '16), that can be viewed as a special case of separation where all particles have the same color. It is significantly more challenging to prove that the desired behavior is achieved in the heterogeneous setting, however, even in the bichromatic case we focus on. This requires combining several new techniques, including the *cluster expansion* from statistical physics, a new variant of the *bridging* argument of Miracle, Pascoe and Randall (RANDOM '11), the *high-temperature expansion* of the Ising model, and careful probabilistic arguments.

2012 ACM Subject Classification Mathematics of computing → Stochastic processes; Theory of computation → Self-organization; Theory of computation → Random walks and Markov chains

Keywords and phrases Markov chains, Programmable matter, Cluster expansion

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2019.54

Category RANDOM

Related Version A full version is available online at `https://arxiv.org/abs/1805.04599`.

Funding *Sarah Cannon*: Supported by National Science Foundation (NSF) award DMS-1803325.
Joshua J. Daymude: Supported by NSF awards CCF-1422603, CCF-1637393, and CCF-1733680.
Cem Gökmen: Supported by NSF award CCF-1733812.
Dana Randall: Supported by NSF awards CCF-1526900, CCF-1637031, and CCF-1733812.
Andréa W. Richa: Supported by NSF awards CCF-1422603, CCF-1637393, and CCF-1733680.

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019).
Editors: Dimitris Achlioptas and László A. Végh; Article No. 54; pp. 54:1–54:22

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Across many disciplines spanning computational, physical, and social sciences, heterogeneous systems self-organize into both separated (or segregated) and integrated states. Examples include molecules exhibiting attractive and repulsive forces, distinct types of bacteria competing for resources while collaborating towards common goals (e.g., [35, 39]), social insects tolerating or aggressing towards those from other colonies (e.g., [30, 20]), and inherent human biases that influence how we form and maintain social groups (e.g., [37, 16]). In each of these, individuals are of different "types": integration occurs when the ensemble gathers together without much preference about the type of their neighbors, while separation occurs when individuals cluster with others of the same type. Here, we investigate these fundamental behaviors of separation or integration as they apply to *programmable matter*, a material that can alter its physical properties based on user input or stimuli from its environment. Instead of studying a particular instantiation of programmable matter, of which there are many [1, 7, 36, 31], we abstractly envision these systems as collections of simple, active computational *particles* that individually execute local distributed algorithms to collectively achieve some emergent behavior. We consider *heterogeneous* particle systems in which particles have immutable *colors*. We seek local, distributed algorithms that, when run by each particle independently and concurrently, result in emergent, self-organizing *separation* or *integration* of color classes.

This work uses the *stochastic approach to self-organizing particle systems* first used for *compression*, where (monochromatic) particles self-organize to gather together as tightly as possible [6]. Using this stochastic approach, one first defines an energy function where desired configurations have the lowest energy values. One then designs a Markov chain whose long run behavior favors these low energy configurations. This Markov chain is carefully designed so that all its transition probabilities can be computed locally, allowing it to be translated to a fully local distributed algorithm each particle can run independently. The resulting collective, emergent behavior of this distributed algorithm is thus described by the long run behavior of the Markov chain. Using this stochastic approach, we previously extended our compression algorithm [6] to an algorithm for *shortcut bridging* [2] — or maintaining bridge structures that balance the tradeoff between bridge efficiency and cost — and developed the theoretical basis for an experimental study in swarm robotics [32]. While the process of designing distributed algorithms for self-organizing particle systems via this stochastic approach is fairly well-understood, proving that such algorithms achieve their desired objectives remains quite challenging. In particular, it is not enough to know the desired configurations have the highest long-run probability; there may be so many other, lower probability configurations that they collectively outweigh the desirable ones. This energy/entropy trade-off has been studied in various Markov chains for the purposes of proving slow mixing, but we analyze it directly to show our algorithms achieve the desired objectives with high probability.

Here, we focus on separation and integration in heterogeneous systems. Our inspiration comes from the classical Ising model in statistical physics [18, 38], where the vertices of a graph are assigned positive and negative "spins" and there are rules governing the probability that adjacent vertices have the same spin. Connected to the Ising model is classical work from stochastic processes on the Schelling model of segregation [33, 34], which explores how individuals' micro-motives can induce macro-level phenomena like racial segregation in residential neighborhoods. Recent variants of this model from computer science have investigated the degree of individual bias required to induce such segregation [5, 17], and a related distributed algorithm has been developed [29]. Our work differs from those on

the Ising and Schelling models because of natural physical constraints on systems of self-organizing, active particles like ours. For example, interpreting particles of one color to be vertices with positive spin and particles of another color to be particles with negative spin, this acts like an Ising model, but on a graph that evolves as particles more. Despite these obstacles, we apply ideas developed for rigorously analyzing the Ising and similar models to prove our distributed algorithm for separation and integration accomplishes the desired goals.

While we are interested in distributed algorithms, it is worth noting that efficient stochastic algorithms for separation can be challenging even with centralized Markov chains. Separation of a region into equitably sized, compact districts has been widely explored recently in the context of gerrymandering, where the aim is to sample colorings of a weighted graph from an appropriately defined stationary distribution [10, 15]. Heuristics for random districting have been discussed in the media, but there are still no known rigorous, efficient algorithms.

1.1 Results

We present a distributed algorithm for self-organizing separation and integration that takes as input two bias parameters, λ and γ. Setting $\lambda > 1$ corresponds to particles favoring having more neighbors; this is known to cause compression in homogeneous systems when λ is large enough [6]. For separation in the heterogeneous setting, we introduce a second parameter γ, where $\gamma > 1$ corresponds to particles favoring having more neighbors *of their own color*. We then investigate for what values of λ and γ our algorithm yields compression and separation. Informally, a particle system is separated if there is a subset of particles such that (i) the boundary between this subset and the rest of the system is small, (ii) a large majority of particles in this subset are of the same color, say c, and (iii) very few particles with color c exist outside of this subset. This notion of separation (defined formally in Definition 3) captures what it means for a system to have large monochromatic regions of particles.

We prove that for any $\lambda > 1$ and $\gamma > 4^{5/4} \sim 5.66$ such that $\lambda\gamma > 2(2+\sqrt{2})e^{0.0003} \sim 6.83$, our algorithm accomplishes separation with high probability.[1] However, we prove the opposite for some values of γ close to one; counterintuitively, this even includes some values of $\gamma > 1$, the regime where particles favor having like-colored neighbors. Formally, we prove that for any $\lambda > 1$ and $\gamma \in (79/81, 81/79)$ such that $\lambda(\gamma+1) > 2(2+\sqrt{2})e^{0.00003} \sim 6.83$, our algorithm fails to achieve separation (i.e., it achieves integration) with high probability.

1.2 Proof Techniques

Because our distributed algorithm is based on a Markov chain, we can use standard tools such as detailed balance to understand its long-term behavior and prove its convergence to a unique probability distribution π over particle system configurations. This stationary distribution π depends on the input parameters λ and γ. Our main contribution is analyzing π for various ranges of λ and γ, showing that a configuration drawn from distribution π is either very likely (for large γ) or very unlikely (for γ close to one) to be separated.

To show separation occurs when λ and γ are both large, we modify the proof technique of *bridging* introduced by Miracle, Pascoe, and Randall [28]. To show separation does not occur when λ is large and γ is small (close to one), we use a probabilistic argument, a Chernoff-type

[1] We say an event A occurs with high probability (w.h.p.) if $\Pr[A] \geq 1 - c^{n^\delta}$, where $0 < c < 1$ and $\delta > 0$ are constants and n is the number of particles. Our w.h.p. results all have $\delta \in \{1/2, 1/2 - \varepsilon\}$, for arbitrarily small $\varepsilon > 0$.

bound, and a decomposition of configurations into different regions. These arguments — both for large and small γ — require that the particle system is compressed; i.e., that the system has perimeter $\Theta(\sqrt{n})$. However, the arguments from [6] showing compression occurs for homogeneous systems when λ is large do not extend to the heterogeneous setting.

We instead turn to the *cluster expansion* from statistical physics to show our separation algorithm achieves compression for large enough γ. The cluster expansion was first introduced in 1937 by Mayer [27], though a more modern treatment can be found in the textbook [12] where it is used to derive several properties of statistical physics models including the Ising and hard-core models. In the past year, the cluster expansion has received renewed attention in the computer science community due to the recent work of Helmuth, Perkins, and Regts that uses the cluster expansion to develop approximate counting and sampling algorithms for low-temperature statistical physics models on lattices including the Potts and hard-core models [14]. Subsequent work has considered similar techniques on expander graphs [19] and random regular bipartite graphs [23]. Inspired by the interpolation method of Barvinok [4, 3], these works give algorithms for estimating partition functions that explicitly calculate the first $\log n$ coefficients of the cluster expansion. We use the cluster expansion differently, to separate the volume and surface contributions to a partition function.

The cluster expansion is a power series representation of $\ln Z$ where Z is a *polymer partition function*. We relate each of our quantities of interest to a particular polymer partition function, and then use a version of the Kotecký-Preiss condition [21] to show that the power series in the cluster expansion is convergent for the ranges of parameters we are interested in. We then use this convergent cluster expansion to split our polymer partition functions into a *volume term*, depending only on the size of the region of interest, and a *surface term*, depending only on its perimeter. This separation into volume and surface terms turns out to be the key to our compression argument, both for large γ and for γ close to one. While splitting partition functions into volume and surface terms is not a new idea in the statistical physics literature (for example, Section 5.7.1 of [12] uses it to derive an explicit expression for the infinite volume pressure of the Ising model on $\mathbb{Z}^d$ with large magnetic field), we are the first to bring this approach into the computer science literature. We are hopeful it will be useful beyond its specific applications in this paper.

2 Background

We begin by defining our amoebot model for programmable matter and stating a few key results. We then extend the amoebot model to heterogeneous particle systems and formally define what it means for a system to be separated or integrated. We conclude with the necessary terminology and results on Markov chains.

2.1 The Amoebot Model

In the *amoebot model*, introduced in [9] and fully described in [8], programmable matter consists of individual, homogeneous computational elements called *particles*. In its geometric variant, particles are assumed to occupy nodes of the triangular lattice $G_\Delta = (V, E)$ and can move along its edges (see Figure 1a). Each node in V can be occupied by at most one particle at a time. Each particle occupies either a single node in V (i.e., it is *contracted*) or a pair of adjacent nodes in V (i.e., it is *expanded*), as in Figure 1b. Particles move via a series of *expansions* and *contractions*: a contracted particle can expand into an unoccupied adjacent node to become expanded, and completes its movement by contracting to once again occupy a single node.

Figure 1 (a) A section of the triangular lattice G_Δ. (b) Expanded and contracted particles (black dots) on G_Δ (gray lattice). Particles with a black line between their nodes are expanded.

Two particles occupying adjacent nodes are said to be *neighbors*. Each particle is *anonymous*, lacking a unique identifier, but can locally identify each of its neighboring locations and can determine which of these are occupied by particles. Each particle has a constant-size local memory that it can write to and its neighbors can read from for communication. In particular, a particle stores whether it is contracted or expanded in its memory. Particles do not have any access to global information such as a shared compass, coordinate system, or estimate of the size of the system.

The system progresses through *atomic actions* according to the standard asynchronous model of computation from distributed computing (see, e.g., [25]). A classical result under this model states that for any concurrent asynchronous execution of atomic actions, there exists a sequential ordering of actions producing the same end result, provided conflicts that arise in the concurrent execution are resolved. In the amoebot model, an atomic action corresponds to the activation of a single particle. Once activated, a particle can (*i*) perform an arbitrary, bounded amount of computation involving information it reads from its local memory and its neighbors' memories, (*ii*) write to its local memory, and (*iii*) perform at most one expansion or contraction. Conflicts involving simultaneous particle expansions into the same unoccupied node are assumed to be resolved arbitrarily such that at most one particle moves to some unoccupied node at any given time. Thus, while in reality many particles may be active concurrently, it suffices when analyzing algorithms under the amoebot model to consider a sequence of activations where only one particle is active at a time.

2.2 Terminology and Results for Homogeneous Particle Systems

We now recall the relevant terminology and notation from our previous work on compression [6]. A particle system *arrangement* is the set of vertices of the triangular lattice G_Δ occupied by particles. Two arrangements are equivalent if they are translations of each other; we define a particle system *configuration* to be an equivalence class of arrangements. An *edge* of a configuration is an edge of G_Δ where both endpoints are occupied by particles. A configuration is *connected* if for any two particles in the system, there is a path of such edges between them. A configuration has a *hole* if there is a maximal, finite, connected component of unoccupied vertices in G_Δ.

As we justify with Lemma 6, our analysis will focus on connected, hole-free configurations. The *boundary* of such a configuration σ is the closed walk $\mathcal{P}$ on edges of σ that encloses all particles of σ and no unoccupied vertices of G_Δ. The *perimeter* $p(\sigma)$ of configuration σ is the length of this walk, also denoted $|\mathcal{P}|$. The following bounds the number of configurations with a given perimeter.

▶ **Lemma 1** ([6], Lemma 4.3). *For any $\nu > 2+\sqrt{2}$, there is an integer $n_1(\nu)$ such that for all $n \geq n_1(\nu)$, the number of connected, hole-free particle system configurations with n particles*

and perimeter k is at most ν^k.

Let $p_{min}(n)$ be the minimum possible perimeter for a configuration of n particles; it is easy to see that $p_{min}(n) = \Theta(\sqrt{n})$. Given any $\alpha > 1$, a configuration of n particles is said to be α-*compressed* if $p(\sigma) \leq \alpha \cdot p_{min}(n)$. The following lemma establishes a concrete upper bound on $p_{min}(n)$.

▶ **Lemma 2.** *For any $n \geq 1$, there is a connected, hole-free particle system configuration of n particles with perimeter at most $2\sqrt{3}\sqrt{n}$. That is, $p_{min}(n) \leq 2\sqrt{3}\sqrt{n}$.*

Proof. This lemma follows easily from noting that hexagonal configurations of n particles have perimeter on the order of $2\sqrt{3}\sqrt{n}$; a proof can be found in Appendix A.1. ◀

2.3 Heterogeneous Particle Systems

Generalizing previous work on the amoebot model in which all particles are homogeneous and indistinguishable, we assume that each particle P has a fixed *color* $c(P) \in \{c_1, \ldots, c_k\}$ that is visible to itself and its neighbors, where $k \ll n$ is a constant. We extend the definition of *configuration* given in Section 2.2 to include both the vertices of G_Δ occupied by particles as well as the colors of those particles. An edge of configuration σ with endpoints occupied by particles P and Q is *homogeneous* if $c(P) = c(Q)$ and *heterogeneous* otherwise.

We further extend the original model by allowing neighboring particles to exchange their positions in a *swap move*. Swap moves have no meaning in homogeneous systems as all particles are indistinguishable, but they grant heterogeneous systems flexibility in allowing particles trapped in the interior of the system to move freely.[2] These swap moves are not necessary for the correctness of our algorithm or our rigorous analysis, but enable faster convergence in practice.

In this paper, we study heterogeneous systems with $k = 2$ color classes. As discussed in Section 5, our algorithm performs well in practice for larger values of k and we expect our proof techniques would generalize without needing significant new ideas. However, this generalization would be cumbersome; thus, for simplicity, we restrict our attention to systems with colors $\{c_1, c_2\}$. For 2-heterogeneous systems, we can formally define separation with respect to having large monochromatic regions.

▶ **Definition 3.** *For $\beta > 0$ and $\delta \in (0, 1/2)$, a 2-heterogeneous particle system configuration σ is said to be* (β, δ)-separated *if there is a subset of particles R such that:*

1. *There are at most $\beta\sqrt{n}$ edges of σ with exactly one endpoint in R;*
2. *The density of particles of color c_1 in R is at least $1 - \delta$; and*
3. *The density of particles of color c_1 not in R is at most δ.*

Unpacking this definition, β controls how small a boundary there is between the monochromatic region R and the rest of the system, with smaller β requiring smaller boundaries. The δ parameter expresses the tolerance for having particles of the wrong color within the monochromatic region R: small values of δ require stricter separation of the color classes, while larger values of δ allow for more integrated configurations. Notably, R does not need to be connected.

[2] In domains where physical swap moves are unrealistic, colors could be treated as in-memory attributes that could be exchanged by neighboring particles to simulate a swap move.

2.4 Markov Chains

A thorough treatment of Markov chains can be found in the standard textbook [22]. A *Markov chain* is a memoryless random process on a state space Ω; for our purposes, Ω is finite and discrete. We focus on discrete time Markov chains, where one transition occurs per *iteration* (or *step*). Because of its stochasticity, we can completely describe a Markov chain by its transition matrix M, which is an $|\Omega| \times |\Omega|$ matrix where for $x, y \in \Omega$, $M(x, y)$ is the probability, if in state x, of transitioning to state y in one step. The t-step transition probability $M^t(x, y)$ is the probability of transitioning from x to y in exactly t steps.

A Markov chain is *ergodic* if it is both *irreducible* (i.e., for all $x, y \in \Omega$ there is a t such that $M^t(x, y) > 0$) and *aperiodic* (i.e., for all $x \in \Omega$, $\gcd\{t : M^t(x, x) > 0\} = 1$) . A *stationary distribution* of a Markov chain is a probability distribution π over Ω such that $\pi M = \pi$. Any finite, ergodic Markov chain converges to a unique stationary distribution given by $\pi(y) = \lim_{t\to\infty} M^t(x, y)$ for any $x, y \in \Omega$; importantly, for such chains this distribution is independent of starting state x. To verify π' is the unique stationary distribution of a finite ergodic Markov chain, it suffices to check that $\pi'(x)M(x, y) = \pi'(y)M(y, x)$ for all $x, y \in \Omega$ (the *detailed balance condition*; see, e.g., [11]).

Given a state space Ω, a set of allowable transitions between states, and a desired stationary distribution π on Ω, the Metropolis-Hastings algorithm [13] gives a Markov chain on Ω with those transitions that converges to π. For separation, the state space contains particle configurations and transitions correspond to configurations that differ by one particle move; the stationary distribution π favors well-separated configurations; and we calculate transition probabilities according to the Metropolis-Hasting algorithm (using a *Metropolis filter*). Importantly, we choose π so that these transition probabilities can be calculated by an individual particle using only information in its local neighborhood.

3 The Separation Algorithm

We now present our stochastic, local, distributed algorithm for separation. Our algorithm achieves separation by biasing particles towards moves that both gain them more neighbors overall and more like-colored neighbors. We use two bias parameters to control this preference: $\lambda > 1$ corresponds to particles favoring having more neighbors, and $\gamma > 1$ corresponds to particles favoring having more neighbors of their own color.

In order to leverage powerful techniques from Markov chain analysis and statistical physics to prove the correctness of our algorithm, we design our algorithm to follow certain invariants. First, assuming the initial particle system configuration is connected, our algorithm ensures it remains connected; this is necessary because particles have strictly local communication abilities so a disconnected particle is unable to communicate with or even find the rest of the particles. Second, our algorithm eventually eliminates all holes in the configuration, and no new holes are ever formed. This is necessary because our proof techniques only apply to hole-free configurations. Third, once all holes have been eliminated, all moves allowed by our algorithm are *reversible*: if a particle moves from node u to an adjacent node v in one step, there is a nonzero probability that it moves back to u in the next step. Finally, the moves allowed by our algorithm suffice to transform any connected, hole-free configuration into any other connected, hole-free configuration.

Our algorithm uses two locally-checkable properties that ensure particles do not disconnect the system or form a hole when moving (our first two invariants). We use the following notation. For a location ℓ — i.e., a node of the triangular lattice G_Δ — let $N_i(\ell)$ denote the particles of color c_i occupying locations adjacent to ℓ. For neighboring locations ℓ and ℓ', let

$N_i(\ell \cup \ell')$ denote the set $N_i(\ell) \cup N_i(\ell')$, excluding particles occupying ℓ and ℓ'. When ignoring color, let $N(\ell) = \bigcup_i N_i(\ell)$; define $N(\ell \cup \ell')$ analogously. Let $\mathbb{S} = N(\ell) \cap N(\ell')$ denote the set of particles adjacent to both locations. A particle can move from location ℓ to ℓ' if one of the following are satisfied:

▷ Property 4. $|\mathbb{S}| \in \{1, 2\}$ and every particle in $N(\ell \cup \ell')$ is connected to exactly one particle in $\mathbb{S}$ by a path through $N(\ell \cup \ell')$.

▷ Property 5. $|\mathbb{S}| = 0$, and both $N(\ell) \setminus \{\ell'\}$ and $N(\ell') \setminus \{\ell\}$ are nonempty and connected.

Note these properties do not need to be verified for swap moves, since swap moves do not change the set of occupied locations and thus cannot disconnect the system or create a hole.

We now define the Markov chain $\mathcal{M}$ for separation. The state space Ω of $\mathcal{M}$ is the set of all connected heterogeneous particle system configurations of n contracted particles, and Algorithm 1 defines its transition probabilities. We note that $\mathcal{M}$, a centralized Markov chain, can be directly translated to a fully distributed, local, asynchronous algorithm $\mathcal{A}$ that can be run by each particle independently and concurrently to achieve the same system behavior. This translation is much the same as for previous algorithms developed using the stochastic approach to self-organizing particle systems [6, 2]; we refer the interested reader to those papers for details. Importantly, this translation is only possible because all probability calculations and property checks in $\mathcal{M}$ use strictly local information available to the particles involved. Simulations of $\mathcal{M}$ can be found in Section 3.2.

Algorithm 1 Markov Chain $\mathcal{M}$ for Separation and Integration

Beginning at any connected configuration σ_0 of n particles, repeat:

Choose a particle P uniformly at random; let c_i be its color and ℓ its location.
Choose a neighboring location ℓ' and $q \in (0, 1)$ each uniformly at random.
if ℓ' is unoccupied **then**
    P expands to occupy both ℓ and ℓ'.
    Let $e = |N(\ell)|$ (resp., $e_i = |N_i(\ell)|$) be the number of neighbors (resp., of color c_i) P had when contracted at location ℓ, and define $e' = |N(\ell')|$ and $e'_i = |N_i(\ell')|$ analogously.
    if (i) $e \neq 5$, (ii) ℓ and ℓ' satisfy Property 4 or 5, and (iii) $q < \lambda^{e'-e} \cdot \gamma^{e'_i - e_i}$ **then**
        P contracts to ℓ'.
    else P contracts back to ℓ.
else if ℓ' is occupied by particle Q of color c_j **then**
    if $q < \gamma^{|N_i(\ell') \setminus \{P\}| - |N_i(\ell)| + |N_j(\ell) \setminus \{Q\}| - |N_j(\ell')|}$ **then** P and Q perform a swap move.

3.1 The Stationary Distribution of Markov Chain $\mathcal{M}$

In this section, we prove that Markov chain $\mathcal{M}$ maintains the four invariants described previously and then characterize its stationary distribution.

▶ **Lemma 6.** *If the particle system is initially connected, it remains connected throughout the execution of $\mathcal{M}$. Moreover, $\mathcal{M}$ eventually eliminates any holes in the initial configuration, after which no holes are ever introduced again.*

Proof. This follows directly from analogous results for compression [6]. Although the separation and compression algorithms assign different probabilities to particle moves, the set of allowed movements is exactly the same, excluding swap moves that do not change the set of occupied nodes of G_Δ, so they cannot disconnect the system or introduce a hole. ◀

▶ **Lemma 7.** *Once all holes have been eliminated, every possible particle move is reversible; that is, if there is a positive probability of moving from configuration σ to configuration τ, then there is a positive probability of moving from τ to σ.*

Proof. Suppose, for example, that a particle P moves from location ℓ to ℓ'. In the next time step, it is possible for P to be chosen again (Step 1) and for ℓ to be chosen as the position to explore (Step 2). Because Properties 4 and 5 are symmetric with respect to ℓ and ℓ', whichever was satisfied in the forward move will also be satisfied in this reverse move. Finally, the probability checked in Condition (*iii*) of Step 7 is always nonzero, so all together there is a nonzero probability that P moves back to ℓ in this reverse move. Swap moves can be shown to be reversible in a similar way. ◀

▶ **Lemma 8.** *Markov chain $\mathcal{M}$ is ergodic on the space of connected, hole-free configurations.*

Proof Sketch. One can show that $\mathcal{M}$ is irreducible (i.e., the moves of $\mathcal{M}$ suffice to transform any configuration to any other configuration) similarly to the proof of the same fact for compression [6]: it is first shown that any configuration can be reconfigured into a straight line; then, the line can be sorted by the color of the particles; finally, by reversibility (Lemma 7), the line can be reconfigured into any configuration. Additionally, it is easy to see that $\mathcal{M}$ is aperiodic: at each iteration of $\mathcal{M}$, there is a nonzero probability that the configuration does not change. Thus, because $\mathcal{M}$ is irreducible and aperiodic, we conclude it is ergodic. ◀

Because $\mathcal{M}$ is finite and ergodic, it converges to a unique stationary distribution π that we now characterize. For a configuration σ, let $h(\sigma)$ be the number of heterogeneous edges in σ.

▶ **Lemma 9.** *For $Z = \sum_\sigma (\lambda\gamma)^{-p(\sigma)} \cdot \gamma^{-h(\sigma)}$, the stationary distribution of $\mathcal{M}$ is:*

$$\pi(\sigma) = \begin{cases} (\lambda\gamma)^{-p(\sigma)} \cdot \gamma^{-h(\sigma)}/Z & \text{if } \sigma \text{ is connected and hole-free;} \\ 0 & \text{otherwise.} \end{cases}$$

Proof Sketch. By Lemma 6, when $\mathcal{M}$ starts at a connected configuration it eventually reaches and remains in the set of configurations that are connected and hole-free. Thus, disconnected configurations and configurations with holes have zero weight at stationarity. In Appendix A.2, we show using detailed balance that the unique stationary distribution of $\mathcal{M}$ can be written, for σ connected and hole-free, as $\pi(\sigma) = \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}/Z_e$ where $e(\sigma)$ is the number of edges and $a(\sigma)$ is the number of homogeneous edges of σ and $Z_e = \sum_\sigma \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}$. This can be rewritten as in the lemma using two facts: (*i*) since every edge is either homogeneous or heterogeneous, $e(\sigma) = a(\sigma) + h(\sigma)$; and (*ii*) for any connected, hole-free configuration σ, $e(\sigma) = 3n - p(\sigma) - 3$, a result shown in [6]. ◀

The remainder of this paper will be spent analyzing this stationary distribution.

3.2 Simulations

We supplement our rigorous results with simulations that show separation occurs for even better values of λ and γ than our proofs guarantee, indicating that our proven bounds are likely not tight. We simulated $\mathcal{M}$ on heterogeneous particle systems with two colors, using 50 particles of each color. Figure 2 shows the progression of $\mathcal{M}$ over time with bias parameters $\lambda = 4$ and $\gamma = 4$, the regime in which particles prefer to have more neighbors, especially those of their own color. The simulation ran for nearly 70 million iterations, but much of the system's compression and separation occurs in the first million iterations. Separation still occurs even when swap moves are disallowed, but takes much longer to achieve.

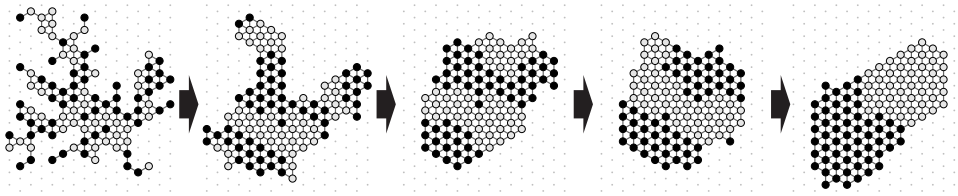

Figure 2 A 2-heterogeneous particle system of 100 particles starting from an arbitrary initial configuration after (from left to right) 0; 50,000; 1,050,000; 17,050,000; and 68,250,000; iterations of $\mathcal{M}$ with $\lambda = 4$ and $\gamma = 4$.

Figure 3 compares the resulting system configurations after running $\mathcal{M}$ from the same initial configuration for the same number of iterations, varying only the values of λ and γ. We observe four distinct phases: compressed-separated, compressed-integrated, expanded-separated, and expanded-integrated. We rigorously verify the compressed-separated behavior (i.e., when λ and γ are large), and do the same for the compressed-integrated behavior (i.e., when λ is large and γ is small). We do not give proofs for expanded configurations; in fact, our current definition of separation may not accurately capture what occurs in expanded configurations.

4 Summary of Results and Proofs

Here we summarize our results and proofs; details have been omitted due to length constraints.

We want to know for which values of λ and γ separation does or does not occur. Our proof techniques only apply to compressed configurations, so we must first show that Markov chain $\mathcal{M}$ achieves compression for the values of λ and γ we are interested in. Previous proofs of compression in homogeneous particle systems break down for heterogeneous systems, so we utilize the *cluster expansion* to overcome this obstacle. The cluster expansion comes from statistical physics and allows us to rewrite a sum over collections of disjoint objects in terms of a sum over collections of overlapping objects. This latter sum is often much easier to work with. For the cluster expansion to be useful, the formal power series it involves must be convergent. We highly recommend Chapter 5 of [12] to learn more about the cluster expansion. Here we present only the relevant definitions and results from this chapter.

In a *polymer model*, we consider a finite set Γ, the elements of which are called *polymers*. We will consider polymers that are collections of edges of G_Δ having certain properties; for large γ, our polymers are minimal cut sets that we call *loops*, and when γ is close to one, our polymers are connected edge sets with an even number of edges incident on each vertex. Formally, polymers only need to satisfy:

- Each polymer $\xi \in \Gamma$ has a real *weight* $w(\xi)$.[3]
- There is a notion of pairwise *compatibility* for polymers.

Polymers are typically compatible when they are well-separated in some sense. Our loop polymers will be compatible when they share no edges, and our even polymers will be compatible when they are not incident on any of the same vertices. We say a collection of polymers $\Gamma' \subseteq \Gamma$ is *compatible* if all polymers in Γ' are pairwise compatible.

[3] In general $w(\xi)$ can be complex, but for our purposes it will always be a (positive or negative) real number.

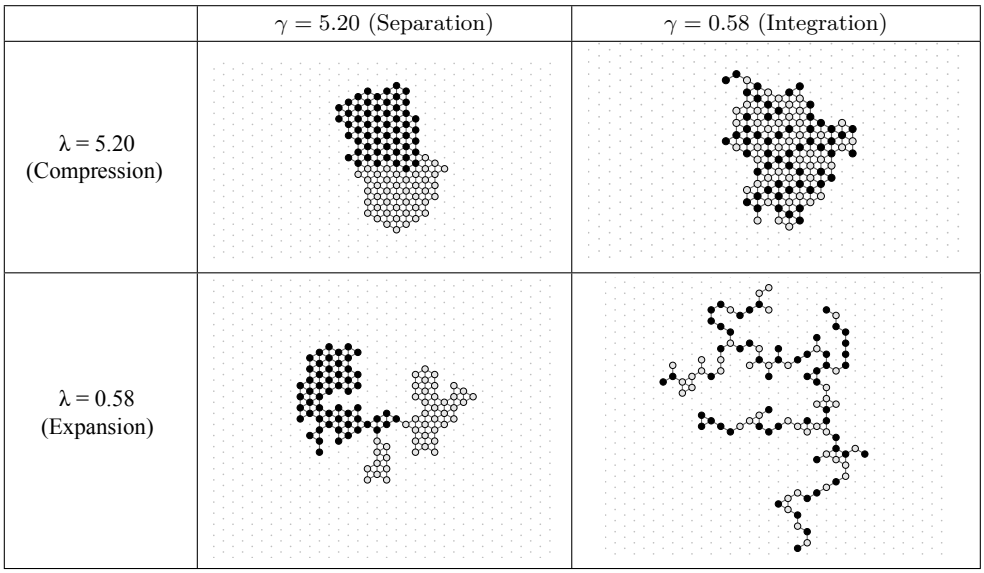

Figure 3 A 2-heterogeneous particle system of 100 particles starting in the leftmost configuration of Figure 2 after 50,000,000 iterations of $\mathcal{M}$ for various values of the parameters λ and γ.

The *polymer partition function* is defined as:

$$\Xi = \sum_{\substack{\Gamma' \subseteq \Gamma \\ compatible}} \prod_{\xi \in \Gamma'} w(\xi).$$

Many partition functions of spin systems, such as the Ising model or the hard-core lattice gas model, can be written in this form as polymer partition functions. Such an abstract sum can sometimes be hard to analyze, but the *cluster expansion* gives a way of rewriting this expression in terms of a sum over subsets $\Gamma' \subseteq \Gamma$ where many polymers are incompatible; because incompatible polymers 'touch,' we can enumerate such collections more easily and thus such sums are often easier to work with

Formally, consider an ordered multiset $X = \{\xi_1, \xi_2, \ldots, \xi_m\} \subseteq \Gamma$. Let H_X be the *incompatibility graph* on vertex set $\{1, 2, \ldots, m\}$ where $i \sim j$ whenever ξ_i and ξ_j are incompatible. We say that the X is a *cluster* if H_X is connected.[4] Let $|X| = m$ denote the number of polymers in cluster X (with polymers counted with the appropriate multiplicities).

The *cluster expansion* is the formal power series for $\ln \Xi$ given in Equation 2. Often this power series does not converge, but the *Kotecky-Preiss condition* guarantees convergence and is often easy to verify [21]. The following theorem states the Kotecky-Preiss condition (Equation 1) and the cluster expansion of Ξ.

▶ **Theorem 10** ([12], Chapter 5). *Let Γ be a finite set of polymers ξ with real weights $w(\xi)$ and a notion of pairwise compatibility. If there exists a function $a : \Gamma \to \mathbb{R}_{>0}$ such that for all $\xi^* \in \Gamma$,*

$$\sum_{\substack{\xi \in \Gamma: \\ \xi, \xi^* \ incompatible}} |w(\xi)| e^{a(\xi)} \leq a(\xi^*), \tag{1}$$

[4] Many sources define clusters to be unordered multisets, necessitating additional combinatorial terms in the cluster expansion; for simplicity, we assume clusters are ordered.

then the polymer partition function Ξ *satisfies*

$$\ln \Xi = \sum_{X\colon cluster} \frac{1}{|X|!} \left(\sum_{\substack{G \subseteq H_X\colon \\ connected, \\ spanning}} (-1)^{|E(G)|} \right) \left(\prod_{\xi \in X} w(\xi) \right), \tag{2}$$

where $G \subseteq H_X$ *means* G *is a subgraph of* H_X*.*

The cluster expansion is derived and this theorem is proved in Chapter 5 of [12], for a slightly different (but equivalent) definition of a cluster.

We apply the cluster expansion twice, with two different notions of polymers and compatibility. In both cases, our polymers will be connected edge sets $\xi \subseteq E(G_\Delta)$, and we use that to state a general result here. Let Γ be an infinite set of such polymers that is invariant under translation and rotation of polymers. Two polymers in Γ will be compatible if they are well-separated in the model-dependent sense described above. Polymers are incompatible when they are 'too close;' for a polymer $\xi \in \Gamma$, let $[\xi] \subseteq E(G_\Delta)$ be the the minimal edge set such that if ξ' is not compatible with ξ, then ξ' must contain an edge of $[\xi]$. We use brackets, consistent with the notation of [12], because this is a type of *closure* of a polymer. For our loop polymers, which are compatible if they share no edges, $[\xi] = \xi$. For our even polymers, which are compatible if they are not incident on any of the same vertices, $[\xi]$ is all edges that share an endpoint with an edges of ξ. We denote the size of this edge set as $|[\xi]|$.

We will be interested in some finite region $\Lambda \subseteq E(G_\Delta)$, and we say $\Gamma_\Lambda \subseteq \Gamma$ is all polymers of Γ whose edges are contained in Λ. Let $\partial\Lambda$ be an edge set such that a cluster containing an edge in Λ and an edge not in Λ must contain an edge of $\partial\Lambda$. We will consider loop polymers with edges from $E_{\mathcal{P}}^{int}$, the set of edges with at least one endpoint strictly inside boundary $\mathcal{P}$, so in this case we use $\Lambda = E_{\mathcal{P}}^{int}$ and $\partial\Lambda$ the edges in $\mathcal{P}$. For even polymers, we use $\Lambda = E_{\mathcal{P}}$, all edges on or inside $\mathcal{P}$, and $\partial\Lambda$ is all edges with one endpoint on $\mathcal{P}$ and the other outside.

The following states the key fact about the cluster expansion that we will need. Namely, when a certain mild condition is satisfied, we can use the cluster expansion to give upper and lower bounds on the polymer partition function for Λ in terms of a volume term, depending only on $|\Lambda|$, and a surface term, depending only on $|\partial\Lambda|$.

▶ **Theorem 11.** *Let* Γ *be an infinite set of polymers* $\xi \subseteq E(G_\Delta)$ *that is closed under translation and rotation, and let* $\Lambda \subseteq E(G_\Delta)$ *be finite. If there is a constant* c *such that for any edge* $e \in E(G_\Delta)$*,*

$$\sum_{\substack{\xi \in \Gamma\colon \\ e \in \xi}} |w(\xi)| e^{c|[\xi]|} \leq c,$$

then for any Λ *the partition function*

$$\Xi_\Lambda := \sum_{\substack{\Gamma' \subseteq \Gamma_\Lambda \\ compatible}} \prod_{\xi \in \Gamma'} w(\xi)$$

satisfies

$$e^{\psi|\Lambda| - c|\partial\Lambda|} \leq \Xi_\Lambda \leq e^{\psi|\Lambda| + c|\partial\Lambda|},$$

for some constant $\psi \in [-c, c]$ *that is independent of* Λ*.*

We prove this theorem in Appendix A.3.

This result is the key step needed to show that when λ and γ are both large, compression occurs; as our techniques for establishing separation first require configurations to be compressed, this is a necessary first step. For compression, we look at the *partition function* $Z_{\mathcal{P}}$ for different fixed boundaries $\mathcal{P}$, where $Z_{\mathcal{P}}$ is the sum over all configurations σ with boundary $\mathcal{P}$ of their weights $(\lambda\gamma)^{-|\mathcal{P}|} \cdot \gamma^{-h(\sigma)}$. We cannot analyze $Z_{\mathcal{P}}$ directly, so we instead relate $Z_{\mathcal{P}}$ to a specific polymer partition function $\Xi^{\mathcal{L}}_{\mathcal{P}}$ which does have a cluster expansion. Using the sufficient condition of Theorem 10, we show the cluster expansion for $\Xi^{\mathcal{L}}_{\mathcal{P}}$ is convergent when $\gamma > 4^{5/4}$. We then use this expression of $\ln \Xi^{\mathcal{L}}_{\mathcal{P}}$ as a convergent power series and Theorem 11 to bound $\Xi^{\mathcal{L}}_{\mathcal{P}}$ in terms of a *volume term*, depending only on the number of particles n, and a *surface term*, depending only on $|\mathcal{P}|$, the length of boundary $\mathcal{P}$.

▶ **Lemma 12.** *When $\gamma > 4^{5/4}$, for $c = 0.0001$, there exists a constant $\psi \in [-c, c]$ that depends on γ but is independent of $\mathcal{P}$ such that*

$$e^{(3n-3)\psi - 3c|\mathcal{P}|} \leq \Xi^{\mathcal{L}}_{\mathcal{P}} \leq e^{(3n-3)\psi + 3c|\mathcal{P}|}.$$

This means, in particular, that the ratios of $\Xi^{\mathcal{L}}_{\mathcal{P}}$ and $\Xi^{\mathcal{L}}_{\mathcal{P}'}$ for different boundaries $\mathcal{P}$ and $\mathcal{P}'$ that enclose the same number n of particles can be bounded by an expression that is exponential in the lengths of these boundaries but independent of n. This is essential to our compression argument, which will focus on boundaries of various lengths. We note that it is straightforward, using the previous lemma, to get similar bounds on $Z_{\mathcal{P}}$, the quantity we are actually interested in. We use this to apply a Peierls argument similar to the one used to show compression in [6]. This argument relates the total weight of undesirable configurations — those with boundaries longer than $\alpha \cdot p_{min}$ for some constant $\alpha > 1$ — to the weight of configurations with minimum perimeter, p_{min}. The result is as follows.

▶ **Theorem 13.** *Consider algorithm $\mathcal{M}$ when there are n total particles of two different colors. For $c = 0.0001$, when constants $\alpha > 1$, $\lambda > 1$, and $\gamma > 4^{5/4}$ satisfy*

$$\frac{2(2+\sqrt{2})e^{3c}}{\lambda\gamma}\left(e^{3c}\lambda\gamma^{3/2}\right)^{1/\alpha} < 1,$$

when n is sufficiently large then for $\mathcal{M}$ with parameters λ and γ, configurations drawn from distribution π are α-compressed with probability at least $1 - \zeta^{\sqrt{n}}$ for some constant $\zeta < 1$.

One corollary is that if $\lambda > 1$ and $\gamma > 4^{5/4}$ such that $\lambda\gamma > 2(2+\sqrt{2})e^{0.0003} \sim 6.83$, there exists a constant α such that a configuration drawn from the stationary distribution π of $\mathcal{M}$ is α-compressed with high probability. (Recall, we say an event A occurs with high probability, or w.h.p., if $\Pr[A] \geq 1 - c^{n^{\delta}}$, where $0 < c < 1$ and $\delta > 0$ are constants. Unless we explicitly state otherwise, it will always be the case that $\delta = 1/2$.) Conversely, for any $\alpha > 1$, there exist λ and γ such that $\mathcal{M}$ with these parameter values achieves α-compression at stationarity w.h.p.

We next show, again when λ and γ are large enough, that separation provably occurs. By the previous theorem, it suffices to show this among compressed configurations. We use a technique known as *bridging* that was developed to analyze molecular mixtures called *colloids* [28]. Adapting the bridging approach to our setting required several new innovations to overcome obstacles such as the irregular shapes of particle system configurations, the non-self-duality of the triangular lattice, the interchangeability between color classes, and other technicalities related to interfaces between particles of different colors. The main result of this section is the following theorem. Recall that for a fixed boundary $\mathcal{P}$, the probability distribution $\pi_{\mathcal{P}}$ is over colored particle configurations with this boundary where $\pi_{\mathcal{P}}(\sigma)$ is proportional to $\gamma^{-h(\sigma)}$.

▶ **Theorem 14.** *Let $\mathcal{P}$ be the boundary of n particles with $|\mathcal{P}| \leq \alpha p_{min}$. For any $\beta > 2\sqrt{3}\alpha$ and any $\delta < 1/2$, if γ is large enough that*

$$3^{\frac{2\alpha\sqrt{3}}{\beta}} 4^{\frac{1+3\delta}{4\delta}} \gamma^{-1+\frac{2\alpha\sqrt{3}}{\beta}} < 1$$

then for sufficiently large n a configuration drawn from $\pi_{\mathcal{P}}$ is (β, δ)-separated with probability at least $1 - \zeta^{\sqrt{n}}$ for some constant $\zeta < 1$.

Combining this with the previous theorem, we see that for any $\lambda > 1$ and $\gamma > 4^{5/4} \sim 5.66$ such that $\lambda\gamma > 2(2+\sqrt{2})e^{0.0003} \sim 6.83$, there exist constants β and δ such that for large enough n, $\mathcal{M}$ provably achieves (β, δ)-separation at stationarity w.h.p. Furthermore, for any $\beta > 2\sqrt{3}$ and any $\delta < 1/2$, there are values for λ and γ such that for large enough n, $\mathcal{M}$ provably achieves (β, δ)-separation at stationarity w.h.p.

We are also able to show that there are some values of γ close to one for which separation does not occur. This counterintuitively includes values where $\gamma > 1$ and particles have a preference for being next to particles of the same color. As we did for large values of γ, we first show that when λ is large and γ is close to one, compression provably occurs. The polymer partition function $\Xi^{\mathcal{L}}_{\mathcal{P}}$ from above does not have a convergent cluster expansion when γ is close to one, so we cannot use it to show compression. Instead, we carefully relate $Z_{\mathcal{P}}$ to a different polymer partition function $\Xi^{HT}_{\mathcal{P}}$ by considering the *high temperature expansion*, which rewrites a sum over configurations with a fixed boundary as a sum over even edge sets within that boundary. The high-temperature expansion is well-studied for the Ising model (see, e.g., [12], Section 3.7.3). We show $\Xi^{HT}_{\mathcal{P}}$ has a convergent cluster expansion when γ is close to one. We then use the cluster expansion for this high temperature representation, much the same as above, to show compression provably occurs.

▶ **Theorem 15.** *Consider algorithm $\mathcal{M}$ when there are n total particles of two different colors. For $a = 10^{-5}$, when constants $\alpha > 1$, $\lambda > 1$, and $\gamma \in (79/81, 81/79)$ satisfy*

$$\frac{2(2+\sqrt{2})e^{3a}}{\lambda(\gamma+1)} \left(\frac{\lambda(\gamma+1)}{2e^{-3a}\left(\frac{79}{81}\right)} \right)^{1/\alpha} < 1$$

when n is sufficiently large then for $\mathcal{M}$ with parameters λ and γ, configurations drawn from $\mathcal{M}$'s stationary distribution π are α-compressed with probability at least $1 - \zeta^{\sqrt{n}}$ for some constant $\zeta < 1$.

This theorem implies that for any $\lambda > 1$ and $\gamma \in (79/81, 81/79)$ such that $\lambda(\gamma+1) > 2(2+\sqrt{2})e^{0.00003} \sim 6.83$, there exists a constant α such that a configuration drawn from the stationary distribution π of $\mathcal{M}$ is α-compressed w.h.p. Conversely, for any $\alpha > 1$ and any $\gamma \in (79/81, 81/79)$, for large enough λ algorithm $\mathcal{M}$ with parameters λ and γ achieves α-compression at stationarity w.h.p.

Once we have shown that compression occurs for large λ and γ near one, we show that among these compressed configurations a large amount of separation between color classes is very unlikely. We prove this with a probabilistic argument in which we find a set of polynomially many events such that if separation occurs, then at least one of these events occurs. We then show that each event occurs with probability at most $\zeta^{n^{1/2-\varepsilon}}$ for some $\zeta < 1$ and arbitrarily small $\varepsilon > 0$, which via a union bound over the polynomial number of events implies separation is very unlikely.

▶ **Theorem 16.** *Let $\mathcal{P}$ be any α-compressed boundary. Let $\delta < 1/4$ and γ close enough to one such that there exists a $\mu \in (\delta/(1-2\delta), 1/2)$ where*

$$\left(\frac{\mu}{1-\mu}\right)^{(\mu-\delta/(1-2\delta))/11} < \gamma < \left(\frac{1-\mu}{\mu}\right)^{(\mu-\delta/(1-2\delta))/11}.$$

For any β and any $c < 1/4$, there is a constant $\zeta < 1$ such that the probability a particle configuration drawn at random from $\pi_{\mathcal{P}}$ is (β, δ)-separated is at most $\zeta^{n^{2c}}$.

Combining this with the results above, we see that for $\lambda > 1$ and $\gamma \in (79/81, 81/79)$ such that $\lambda(\gamma+1) > 2(2+\sqrt{2})e^{0.00003} \sim 6.83$, there are constants β and δ such that the probability $\mathcal{M}$ with parameters λ and γ achieves (β, δ)-separation at stationarity is at most $\zeta^{n^{1/2-\varepsilon}}$, where $\varepsilon > 0$ and $\zeta < 1$. Conversely, for any $\beta > 0$ and any $\delta < 1/4$, there exists λ and γ such that $\mathcal{M}$ with these parameters achieves (β, δ)-separation at stationarity with probability at most $\zeta^{n^{1/2-\varepsilon}}$ for $\varepsilon > 0$ and $\zeta < 1$.

5 Conclusion

We considered separation with two colors, but expect our proofs to generalize in a straightforward way to heterogeneous systems with more colors using insights that generalize cluster expansion polymers from the Ising model to the Potts model (see the notion of a *contour* in Pirogov-Sinai theory, e.g., in Chapter 7 of [12]). The proofs would follow the same strategy for two colors, requiring little additional insight but a fairly large amount of technical detail.

We note that, as with previous papers using stochastic, distributed algorithms for programmable matter, we are unable to give any nontrivial bounds on the mixing time of our Markov chain $\mathcal{M}$. The difficulties in proving polynomial upper bounds on the mixing time are unsurprising, given similarities between $\mathcal{M}$ and a well-studied open problem in statistical physics about the mixing time of Glauber dynamics of the Ising model on $\mathbb{Z}^2$ with plus boundary conditions starting from the all minus state [26, 24] (see remarks concluding [6]). However, the mixing time may not be the best bound for characterizing when compression and separation occur. Simulations show that both compression and separation occur fairly quickly (Figure 2), although the algorithm continues to gradually achieve more compression and separation, confirming we likely achieve these goals well before converging to stationarity.

We believe the stochastic approach to self-organizing particle systems, used here to develop a distributed algorithm for separation and integration in programmable matter, is much more broadly applicable. This approach can potentially be applied to any objective described by a global energy function (where the desirable configurations have low energy values), provided changes in energy due to particle movements can be calculated with only local information. Choosing the correct global energy function is the key; translating the energy function into a Markov chain and then into a distributed algorithm is, by now, fairly routine (see [6, 2]). However, proving that the stationary distribution has our desired properties with high probability remains challenging, requiring application-specific proof techniques.

Last, we believe the proof techniques developed here extend beyond our current work. For separation and integration, the key ingredient is the cluster expansion, used recently to develop efficient low-temperature approximations and sampling algorithms, and the related Pirogov-Sinai theory, used to show slow mixing of certain Markov chains. Here, however, we used a completely different aspect of the cluster expansion by separating partition functions into surface and volume terms. The cluster expansion and Pirogov-Sinai theory have been widely used in statistical physics for many purposes, and we believe there are many more ways a thorough understanding of these methods can benefit computer science.

References

1. Leonard M. Adleman. Molecular computation of solutions to combinatorial problems. *Science*, 266(5187):1021–1024, 1994.
2. Marta Andrés Arroyo, Sarah Cannon, Joshua J. Daymude, Dana Randall, and Andréa W. Richa. A stochastic approach to shortcut bridging in programmable matter. *Natural Computing*, 17(4):723–741, 2018.
3. Alexander I. Barvinok. *Combinatorics and complexity of partition functions*, volume 30 of *Algorithms and Combinatorics*. Springer International Publishing, 2016.
4. Alexander I. Barvinok and Pablo Soberón. Computing the partition function for graph homomorphisms with multiplicities. *Journal of Combinatorial Theory, Series A*, 137:1–26, 2016.
5. Prateek Bhakta, Sarah Miracle, and Dana Randall. Clustering and mixing times for segregation models on $\mathbb{Z}^2$. In *Proceedings of the Twenty-fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 327–340, 2014.
6. Sarah Cannon, Joshua J. Daymude, Dana Randall, and Andréa 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, pages 279–288, Chicago, IL, USA, 2016. ACM. A significantly updated version is available at `https://arxiv.org/abs/1603.07991`.
7. David Correa, Athina Papadopoulou, Christophe Guberan, Nynika Jhaveri, Steffen Reichert, Achim Menges, and Skylar Tibbits. 3D-printed wood: Programming hygroscopic material transformations. *3D Printing and Additive Manufacturing*, 2(3):106–116, 2015.
8. Joshua J. Daymude, Kristian Hinnenthal, Andréa W. Richa, and Christian Scheideler. Computing by programmable particles. In *Distributed Computing by Mobile Entities: Current Research in Moving and Computing*, pages 615–681. Springer, Cham, 2019.
9. Zahra Derakhshandeh, Shlomi Dolev, Robert Gmyr, Andréa W. Richa, Christian Scheideler, and Thim Strothmann. Brief announcement: amoebot - a new model for programmable matter. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '14, pages 220–222, New York, NY, USA, 2014. ACM.
10. Moon Duchin and Bridget E. Tenner. Discrete geometry for electoral geography. Preprint available online at `https://arxiv.org/abs/1808.05860`, 2018.
11. William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, New York, 1968.
12. Sacha Friedli and Yvan Velenik. *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press, Cambridge, 2018.
13. Wilfred K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
14. Tyler Helmuth, Will Perkins, and Guus Regts. Algorithmic Pirogov-Sinai theory. In *Proceedings of the 51st ACM Symposium on Theory of Computing*, STOC '19. ACM, 2019.
15. Gregory Herschlag, Han Sung Kang, Justin Luo, Christy V. Graves, Sachet Bangia, Robert Ravier, and Jonathan C. Mattingly. Quantifying gerrymandering in North Carolina. Preprint available online at `https://arxiv.org/abs/1801.03783`, 2018.
16. Michael A. Hogg and John C. Turner. Interpersonal attraction, social identification and psychological group formation. *European Journal of Social Psychology*, 15(1):51–66, 1985.
17. Nicole Immorlica, Robert Kleinberg, Brendan Lucier, and Morteza Zadomighaddam. Exponential segregation in a two-dimensional Schelling model with tolerant individuals. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 984–993, 2017.
18. Ernst Ising. Beitrag zur theorie des ferromagnetismus [contribution to the theory of ferromagnetism]. *Zeitschrift für Physik*, 31(1):253–258, 1925.

19 Matthew Jenssen, Peter Keevash, and Will Perkins. Algorithms for #BIS-hard problems on expander graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 2235–2247, 2019.

20 Brian R. Johnson, Ellen van Wilgenburg, and Neil D. Tsutsui. Nestmate recognition in social insects: overcoming physiological constraints with collective decision making. *Behavioral Ecology and Sociobiology*, 65(5):935–944, 2011.

21 Roman Kotecký and David Preiss. Cluster expansion for abstract polymer models. *Communications in Mathematical Physics*, 103:491–498, 1986.

22 David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, USA, 2009.

23 Chao Liao, Jiabao Lin, Pinyan Lu, and Zhenyu Mao. Counting independent sets and colorings on random regular bipartite graphs. 2019.

24 Eyal Lubetzky, Fabio Martinelli, Alan Sly, and Fabio Lucio Toninelli. Quasi-polynomial mixing of the 2D stochastic Ising model with "plus" boundary up to criticality. *Journal of the European Mathematical Society (JEMS)*, 15(2):339—-386, 2013.

25 Nancy Lynch. *Distributed Algorithms*. Morgan Kauffman, San Francisco, CA, USA, 1996.

26 Fabio Martinelli and Fabio Lucio Toninelli. On the mixing time of the 2D stochastic Ising model with "plus" boundary conditions at low temperature. *Communications in Mathematical Physics*, 296(1):175–213, 2010.

27 Joseph E. Mayer. The statistical mechanics of condensing systems. I. *The Journal of Chemical Physics*, 5:67–73, 1937.

28 Sarah Miracle, Dana Randall, and Amanda Pascoe Streib. Clustering in interfering binary mixtures. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, APPROX/RANDOM 2011, pages 652–663, 2011.

29 Hamed Omidvar and Massimo Franceschetti. Self-organized segregation on the grid. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, pages 401–410, New York, NY, USA, 2017. ACM.

30 T'ai H. Roulston, Grzegorz Buczkowski, and Jules Silverman. Nestmate discrimination in ants: effect of bioassay on aggressive behavior. *Insectes Sociaux*, 50(2):151–159, 2003.

31 Michael Rubenstein, Alejandro Cornejo, and Radhika Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198):795–799, 2014.

32 William Savoie, Sarah Cannon, Joshua J. Daymude, Ross Warkentin, Shengkai Li, Andréa W. Richa, Dana Randall, and Daniel I. Goldman. Phototactic supersmarticles. *Artificial Life and Robotics*, 23(4):459–468, 2018.

33 Thomas C. Schelling. Models of segregation. *The American Economic Review*, 59(2):488–493, 1969.

34 Thomas C. Schelling. Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2):143–186, 1971.

35 Philip S. Stewart and Michael J. Franklin. Physiological heterogeneity in biofilms. *Nature Reviews Microbiology*, 6:199–210, 2008.

36 Rohan Thakker, Ajinkya Kamat, Sachin Bharambe, Shital Chiddarwar, and K. M. Bhurchandi. ReBiS - reconfigurable bipedal snake robot. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 309–314, 2014.

37 John C. Turner. Towards a cognitive redefinition of the social group. *Cahiers de Psychologie Cognitive/Current Psychology of Cognition*, 1(2):93–118, 1981.

38 Dejan Vinković and Alan Kirman. A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103(51):19261–19265, 2006.

39 Guopeng Wei, Connor Walsh, Irina Cazan, and Radu Marculescu. Molecular tweeting: Unveiling the social network behind heterogeneous bacteria populations. In *Proceedings of the 6th ACM Conference on Bioinformatics, Computational Biology and Health Informatics*, BCB '15, pages 366–375, New York, NY, USA, 2015. ACM.

A Appendix

Here we include the proofs of some of our claims that were omitted from the main body of this paper for conciseness and clarity. We do not include any detailed proofs of our technical results due to length constraints.

A.1 Proof of Lemma 2

Recall that Lemma 2 states that for any $n \geq 1$, there is a connected, hole-free particle configuration of n particles with perimeter at most $2\sqrt{3}\sqrt{n}$. That is, $p_{min}(n) \leq 2\sqrt{3}\sqrt{n}$.

Proof. The lemma can easily be verified for $n \leq 6$. For $n \geq 7$, we begin with the case where $n = 3\ell^2 + 3\ell + 1$ for some integer $\ell \geq 1$. A regular hexagon with side length ℓ can be decomposed into six triangles, each with $\ell(\ell+1)/2$ particles, and a single center vertex, for $3\ell^2 + 3\ell + 1$ total particles; see Figure 4a. Such a hexagon has perimeter 6ℓ. We see that

$$p_{min}(3\ell^2 + 3\ell + 1) \leq 6\ell \leq 2\sqrt{3}\sqrt{3\ell(\ell+1)} \leq 2\sqrt{3}\sqrt{n-1} \leq 2\sqrt{3}\sqrt{n}.$$

Now we consider $n = 3\ell^2 + 3\ell + 1 + k$, for integers ℓ and k, where $k \in [1, 6\ell + 6)$. As $(3\ell^2 + 3\ell + 1) + 6\ell + 6 = 3(\ell+1)^2 + 3(\ell+1) + 1$, this covers all possible values of n. We construct a particle configuration on $n = 3\ell^2 + 3\ell + 1 + k$ particles by first constructing a regular hexagon of side length ℓ and then adding the remaining k particles around the outside of this hexagon in a single layer, completing one side before beginning the next; see Figure 4b, where $\ell = 3$ and $k = 6$. For $k \leq \ell$, the perimeter of this configuration is $6\ell + 1$. More generally, the perimeter increases by one when particles begin to be added to a new side of the hexagon, and so for $i = 2, 3, 4, 5, 6$, for $(i-1)\ell + (i-2) < k \leq i\ell + (i-1)$ the perimeter of this configuration is $6\ell + i$. We see that (using $i \leq 6$ and $\ell \geq 1$), for any $i = 1, 2, 3, 4, 5, 6$,

$$\begin{aligned} p_{min}(3\ell^2 + 3\ell + 1 + k) &\leq 6\ell + i \leq 2\sqrt{3}\sqrt{\left(\sqrt{3}\ell + \frac{i}{2\sqrt{3}}\right)^2} = 2\sqrt{3}\sqrt{3\ell^2 + \frac{i^2}{12} + i} \\ &\leq 2\sqrt{3}\sqrt{3\ell^2 + 3 + i} \\ &\leq 2\sqrt{3}\sqrt{3\ell^2 + 3\ell + 1 + i - 1} \\ &\leq 2\sqrt{3}\sqrt{3\ell^2 + 3\ell + 1 + k} = 2\sqrt{3}\sqrt{n}. \end{aligned}$$

This concludes our proof. ◀

(a) (b)

Figure 4 (a) The regular hexagon with side length $\ell = 3$ with $3\ell^2 + 3\ell + 1$ total particles. (b) A configuration with $n = 3\ell^2 + 3\ell + 1 + k$ particles for $\ell = 3$ and $k = 6$ with perimeter $20 < 2\sqrt{3}\sqrt{n}$.

A.2 Detailed Balance Proof that π is the Stationary Distribution of $\mathcal{M}$

Recall that Lemma 9 states that the stationary distribution of $\mathcal{M}$ is given by $\pi(\sigma) = 0$ if σ is disconnected or has holes, and by $\pi(\sigma) = (\lambda\gamma)^{-p(\sigma)} \cdot \gamma^{-h(\sigma)}/Z$ otherwise, where $Z = \sum_\sigma (\lambda\gamma)^{-p(\sigma)} \cdot \gamma^{-h(\sigma)}$. Here, we analyze the necessary cases to verify this with detailed balance.

Proof. We first verify that $\pi(\sigma) = \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}/Z_e$ — where $e(\sigma)$ is the number of edges of σ, $a(\sigma)$ is the number of homogeneous edges of σ, and $Z_e = \sum_\sigma \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}$ — is the stationary distribution by detailed balance. We then show that this form of π can be rewritten as in the lemma.

Consider any two connected, hole-free configurations σ, τ that differ by one move of some particle from location ℓ in σ to a neighboring location ℓ' in τ. By examining $\mathcal{M}$, we see that the probability of transitioning from σ to τ is:

$$M(\sigma,\tau) = \min\left\{1, \lambda^{|N(\ell')|-|N(\ell)|} \cdot \gamma^{|N_i(\ell')|-|N_i(\ell)|}\right\}/6n.$$

A similar analysis shows:

$$M(\tau,\sigma) = \min\left\{1, \lambda^{|N(\ell)|-|N(\ell')|} \cdot \gamma^{|N_i(\ell)|-|N_i(\ell')|}\right\}/6n.$$

Without loss of generality, suppose $\lambda^{|N(\ell')|-|N(\ell)|} \cdot \gamma^{|N_i(\ell')|-|N_i(\ell)|} < 1$, meaning $M(\sigma,\tau)$ is this value over $6n$ and $M(\tau,\sigma) = 1/6n$. Because the only edges that differ in σ and τ are incident to ℓ or ℓ',

$$\begin{aligned}
\pi(\sigma)M(\sigma,\tau) &= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{Z_e} \cdot \frac{1}{n} \cdot \frac{1}{6} \cdot \lambda^{|N(\ell')|-|N(\ell)|} \cdot \gamma^{|N_i(\ell')|-|N_i(\ell)|} \\
&= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{Z_e} \cdot \frac{1}{n} \cdot \frac{1}{6} \cdot \lambda^{e(\tau)-e(\sigma)} \cdot \gamma^{a(\tau)-a(\sigma)} \\
&= \frac{\lambda^{e(\tau)} \cdot \gamma^{a(\tau)}}{Z_e} \cdot \frac{1}{n} \cdot \frac{1}{6} \cdot 1 = \pi(\tau)M(\tau,\sigma)
\end{aligned}$$

Thus, detailed balance is satisfied for particle moves that are not swaps.

Suppose instead that σ and τ differ by a swap move of particle P with color c_i at location ℓ in σ and particle Q with color c_j at neighboring location ℓ' in σ. This move could occur if P or Q is chosen in Step 1 of $\mathcal{M}$, so:

$$M(\sigma,\tau) = \min\left\{1, \gamma^{|N_i(\ell')\setminus\{P\}|-|N_i(\ell)|+|N_j(\ell)\setminus\{Q\}|-|N_j(\ell')|}\right\}/3n.$$

Similarly, because τ has P at location ℓ' and Q at location ℓ, we have:

$$M(\tau,\sigma) = \min\left\{1, \gamma^{|N_i(\ell)\setminus\{P\}|-|N_i(\ell')|+|N_j(\ell')\setminus\{Q\}|-|N_j(\ell)|}\right\}/3n.$$

Without loss of generality, suppose that $\gamma^{|N_i(\ell')\setminus\{P\}|-|N_i(\ell)|+|N_j(\ell)\setminus\{Q\}|-|N_j(\ell')|} < 1$, so $M(\sigma,\tau)$ is this value over $3n$ and $M(\tau,\sigma) = 1/3n$. Then,

$$\begin{aligned}
\pi(\sigma)M(\sigma,\tau) &= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{Z_e} \cdot \frac{2}{n} \cdot \frac{1}{6} \cdot \gamma^{|N_i(\ell')\setminus\{P\}|-|N_i(\ell)|+|N_j(\ell)\setminus\{Q\}|-|N_j(\ell')|} \\
&= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{Z_e} \cdot \frac{2}{n} \cdot \frac{1}{6} \cdot \gamma^{(|N_i(\ell')\setminus\{P\}|+|N_j(\ell)\setminus\{Q\}|)-(|N_i(\ell)|+|N_j(\ell')|)} \\
&= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{Z_e} \cdot \frac{2}{n} \cdot \frac{1}{6} \cdot \gamma^{a(\tau)-a(\sigma)} \\
&= \frac{\lambda^{e(\tau)} \cdot \gamma^{a(\tau)}}{Z_e} \cdot \frac{2}{n} \cdot \frac{1}{6} \cdot 1 = \pi(\tau)M(\tau,\sigma)
\end{aligned}$$

In both cases, detailed balance is satisfied, so we conclude the stationary distribution π (which is only non-zero over connected, hole-free configurations) is given by $\pi(\sigma) = \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}/Z_e$.

Since every edge of σ is either homogeneous or heterogeneous, we have $e(\sigma) = a(\sigma) + h(\sigma)$. From [6], we have $e(\sigma) = 3n - p(\sigma) - 3$, where n is the number of particles in the system. Thus, we can rewrite this unique stationary distribution as follows:

$$\begin{aligned}
\pi(\sigma) &= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{Z_e} \\
&= \frac{\lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}}{\sum_\sigma \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}} \\
&= \frac{(\lambda\gamma)^{-3n+3} \cdot (\lambda\gamma)^{e(\sigma)} \cdot \gamma^{a(\sigma)-e(\sigma)}}{(\lambda\gamma)^{-3n+3} \cdot \sum_\sigma (\lambda\gamma)^{e(\sigma)} \cdot \gamma^{a(\sigma)-e(\sigma)}} \\
&= \frac{(\lambda\gamma)^{e(\sigma)-3n+3} \cdot \gamma^{a(\sigma)-e(\sigma)}}{\sum_\sigma (\lambda\gamma)^{e(\sigma)-3n+3} \cdot \gamma^{a(\sigma)-e(\sigma)}} \\
&= \frac{(\lambda\gamma)^{-p(\sigma)} \cdot \gamma^{-h(\sigma)}}{\sum_\sigma (\lambda\gamma)^{-p(\sigma)} \cdot \gamma^{-h(\sigma)}}.
\end{aligned}$$

This concludes our proof. ◀

A.3 Proof of Boundary-Volume Decomposition of Cluster Expansion

In this section we provide the proof of Theorem 11, which is our decomposition of a polymer partition function into boundary and volume terms via the cluster expansion. For the sake of clarity we restate this theorem here, including all of its hypotheses and assumptions.

▶ **Theorem 11.** *Let Γ be an infinite set of polymers $\xi \subseteq E(G_\Delta)$ that is closed under translation and rotation, and let $\Lambda \subseteq E(G_\Delta)$ be finite. If there is a constant c such that for any edge $e \in E(G_\Delta)$,*

$$\sum_{\substack{\xi \in \Gamma: \\ e \in \xi}} |w(\xi)| e^{c|[\xi]|} \leq c, \tag{3}$$

then for any Λ the partition function

$$\Xi_\Lambda := \sum_{\substack{\Gamma' \subseteq \Gamma_\Lambda \\ compatible}} \prod_{\xi \in \Gamma'} w(\xi)$$

satisfies

$$e^{\psi|\Lambda| - c|\partial\Lambda|} \leq \Xi_\Lambda \leq e^{\psi|\Lambda| + c|\partial\Lambda|},$$

for some constant $\psi \in [-c, c]$ that is independent of Λ.

Proof. We follow the same outline as the proof of the same fact for the Ising model in Section 5.7.1 of [12].

Let $\mathcal{X}$ be all clusters comprised of polymers from Γ, and let $\mathcal{X}_\Lambda$ be all clusters of polymers in Γ_Λ. Note that Equation 3 implies the hypothesis of Theorem 10 (Equation 1) is satisfied, with function $a : \Gamma \to \mathbb{R}$ given by $a(\xi) = c|[\xi]|$:

$$\sum_{\substack{\xi \in \Gamma: \\ \xi, \xi^* \ incompatible}} |w(\xi)| e^{a(\xi)} \leq \sum_{e \in [\xi^*]} \sum_{\substack{\xi \in \Gamma: \\ e \in \xi}} |w(\xi)| e^{c|[\xi]|} \leq c|[\xi^*]|.$$

Because this hypothesis is satisfied for all $\xi^* \in \Gamma$, it certainly holds when we restrict our attention to polymers in Γ_Λ. By Theorem 10, because Γ_Λ is a finite set, this means the cluster expansion for Ξ_Λ converges:

$$\ln \Xi_\Lambda = \sum_{X \in \mathcal{X}_\Lambda} \Psi(X)$$

Let $\overline{X} = \cup_{\xi \in X} \xi$ be the *support* of cluster X and $|\overline{X}|$ the size of this support. Using Equation 3 and standard techniques (see [12], the proof of Theorem 5.4 and Equation (5.29)), the translation and rotation invariance of Γ imply that for any edge $e \in E(G_\Delta)$,

$$\sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}}} |\Psi(X)| \leq c. \tag{4}$$

The proof of this fact is the reason we need a slightly stronger hypothesis (Equation 3) than is needed to guarantee the cluster expansion converges (Equation 1).

For any cluster $X \in \mathcal{X}_\Lambda$, it trivially holds that $1 = (\sum_{e \in \Lambda} \mathbf{1}_{e \in \overline{X}})/\overline{X}$. We can use this fact to rewrite the cluster expansion for Ξ_Λ:

$$\begin{aligned}
\ln \Xi_\Lambda &= \sum_{X \in \mathcal{X}_\Lambda} \Psi(X) = \sum_{\substack{X \in \mathcal{X}: \\ \overline{X} \subseteq \Lambda}} \Psi(X) = \sum_{e \in \Lambda} \sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}, \\ \overline{X} \subseteq \Lambda}} \frac{1}{|\overline{X}|} \Psi(X) \\
&= \sum_{e \in \Lambda} \left(\sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}}} \frac{1}{|\overline{X}|} \Psi(X) - \sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}, \\ \overline{X} \not\subseteq \Lambda}} \frac{1}{|\overline{X}|} \Psi(X) \right) \\
&= \left(\sum_{e \in \Lambda} \sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}}} \frac{1}{|\overline{X}|} \Psi(X) \right) - \left(\sum_{e \in \Lambda} \sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}, \\ \overline{X} \not\subseteq \Lambda}} \frac{1}{|\overline{X}|} \Psi(X) \right).
\end{aligned} \tag{5}$$

The two infinite sums in parentheses above are absolutely convergent by Equation 4, so this difference is well-defined.

To analyze the first term of Equation 5, we note that by the translation and rotation invariance of Γ, the sum

$$\psi := \sum_{\substack{X \in \mathcal{X}: \\ e \in \overline{X}}} \frac{1}{|\overline{X}|} \Psi(X)$$

is independent of e and of Λ and only depends on our particular polymer model; this is the value ψ that appears in the statement of the theorem, and by Equation 4, $|\psi| \leq c$. We conclude the first term of Equation 5 is $\psi|\Lambda|$.

To analyze the second term of Equation 5, recall if cluster X satisfies both $e \in \overline{X}$ for some $e \in \Lambda$ and $\overline{X} \not\subseteq \Lambda$, then $\overline{X}$ must contain some edge $f \in \partial\Lambda$. We rewrite the absolute

value of this second sum as

$$\left|\sum_{e\in\Lambda}\sum_{\substack{X\in\mathcal{X}:\\ e\in\overline{X},\\ \overline{X}\not\subseteq\Lambda}}\frac{1}{|\overline{X}|}\Psi(X)\right| \le \sum_{e\in\Lambda}\sum_{\substack{X\in\mathcal{X}:\\ e\in\overline{X},\\ \overline{X}\not\subseteq\Lambda}}\frac{1}{|\overline{X}|}\,|\Psi(X)|$$

$$\le \sum_{f\in\partial\Lambda}\sum_{\substack{X\in\mathcal{X}:\\ f\in\overline{X}}}|\overline{X}\cap\Lambda|\frac{1}{|\overline{X}|}\,|\Psi(X)|$$

$$\le \sum_{f\in\partial\Lambda}\sum_{\substack{X\in\mathcal{X}:\\ f\in\overline{X}}}|\Psi(X)| \le c\,|\partial\Lambda|\,.$$

The last inequality above follows from Equation 4 and the translation and rotation invariance of Λ.

We conclude that Equation 5 implies

$$\psi|\Lambda| - c|\partial\Lambda| \le \ln\Xi_\Lambda \le \psi|\Lambda| + c|\partial\Lambda|.$$

Exponentiation proves the theorem. ◀