

Local Stochastic Algorithms for Alignment in Self-Organizing Particle Systems

Hridesh Kedia ✉

Georgia Institute of Technology, Atlanta, GA, USA

Shunhao Oh ✉

Georgia Institute of Technology, Atlanta, GA, USA

Dana Randall ✉ 

Georgia Institute of Technology, Atlanta, GA, USA

Abstract

We present local distributed, stochastic algorithms for *alignment* in self-organizing particle systems (SOPS) on two-dimensional lattices, where particles occupy unique sites on the lattice, and particles can make spatial moves to neighboring sites if they are unoccupied. Such models are abstractions of programmable matter, composed of individual computational particles with limited memory, strictly local communication abilities, and modest computational capabilities. We consider oriented particle systems, where particles are assigned a vector pointing in one of q directions, and each particle can compute the angle between its direction and the direction of any neighboring particle, although without knowledge of global orientation with respect to a fixed underlying coordinate system. Particles move stochastically, with each particle able to either modify its direction or make a local spatial move along a lattice edge during a move. We consider two settings: (a) where particle configurations must remain simply connected at all times and (b) where spatial moves are unconstrained and configurations can disconnect.

Our algorithms are inspired by the *Potts model* and its planar oriented variant known as the *planar Potts model* or *clock model* from statistical physics. We prove that for any $q \geq 2$, by adjusting a single parameter, these self-organizing particle systems can be made to collectively align along a single dominant direction (analogous to a solid or ordered state) or remain non-aligned, in which case the fraction of particles oriented along any direction is nearly equal (analogous to a gaseous or disordered state). In the connected SOPS setting, we allow for two distinct parameters, one controlling the ferromagnetic attraction between neighboring particles (regardless of orientation) and the other controlling the preference of neighboring particles to align. We show that with appropriate settings of the input parameters, we can achieve *compression* and *expansion*, controlling how tightly gathered the particles are, as well as *alignment* or *nonalignment*, producing a single dominant orientation or not. While alignment is known for the Potts and clock models at sufficiently low temperatures, our proof in the SOPS setting are significantly more challenging because the particles make spatial moves, not all sites are occupied, and the total number of particles is fixed.

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

Keywords and phrases Self-organizing particle systems, alignment, Markov chains, active matter

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2022.14

Category RANDOM

Related Version *Full Version*: <https://arxiv.org/abs/2207.07956>

Funding *Hridesh Kedia*: Supported by ARO MURI award W911NF-19-1-0233.

Shunhao Oh: Supported by NSF award CCF-1733812 and ARO MURI award W911NF-19-1-0233.

Dana Randall: Supported by NSF awards CCF-1733812 and CCF-2106687 and ARO MURI award W911NF-19-1-0233.



© Hridesh Kedia, Shunhao Oh, and Dana Randall;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022).

Editors: Amit Chakrabarti and Chaitanya Swamy; Article No. 14; pp. 14:1–14:20



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Autonomous, locally interacting agents can collectively organize to accomplish a variety of complex tasks such as foraging for food, building large-scale structures, and transporting objects many times heavier than their weight, as is routinely observed in the living world, in swarms of ants, flocks of birds, and schools of fish [34, 39, 38, 33]. A key component of these diverse self-organized behaviors is achieving consensus in large collectives of autonomous agents with only local interactions. The problem of achieving alignment in collectives of directed agents is an important example of such a consensus problem, and is a fundamental aspect of *flocking*: large scale collective motion in swarms of motile agents [34, 37, 28, 41, 38, 1]. While flocking has been studied extensively [18, 28, 36, 1] with few rigorous results, the more basic problem of alignment has received considerably less attention.

Here, we study alignment in self-organizing particle systems (SOPS) – a collection of simple, active computational particles that individually execute local distributed algorithms. We consider *oriented particle systems* on a two-dimensional lattice, where particles are oriented in one of q directions (with no global compass), for $q \geq 2$, and at most one particle occupies each lattice site. Particles perform moves independently and concurrently by making spatial moves to neighboring empty sites or reorient themselves in new directions with the goal of reaching nearly global alignment.

We consider a stochastic approach, used previously in [7, 8] to achieve *compression*, where connected sets of homogeneous particles self-organize to gather together tightly, *separation* in heterogeneous particle systems, where all of the particles compress, but also gather most tightly with other particles of the same type [5, 6], and *aggregation* of homogeneous particles that are not required to be connected, where most particles accumulate in a small, compact neighborhood [21]. In all of these, phase changes were used to characterize desirable behaviors at stationarity, with high probability. Following a similar approach, we begin by defining an energy function that assigns the highest weight (or lowest energy) to preferable configurations, and design a Markov chain whose long term behavior favors these low energy configurations using transition probabilities given by the Metropolis-Hastings algorithm [25, 15]. We ensure that the transition probabilities of the Markov chain can be computed locally and asynchronously, allowing them to be easily translated to a fully local, distributed algorithm that each particle can run independently. The collective behavior of this distributed algorithm is thus described by the long term behavior of the Markov chain.

1.1 Related work

The alignment problems we study can be viewed as finite, unsaturated variants of the ferromagnetic *Potts model* from statistical physics [40], and a related model known as the *clock* or *planar Potts model* [40, 29]. In the Potts model, vertices of a graph G are assigned one of q possible “spins,” represented here as orientations, and neighboring sites prefer to agree. Let $J > 0$ be a parameter related to inverse temperature and let $\delta(X, Y) = 1$ if $X = Y$ and 0 otherwise. Then the probability of a standard Potts configuration σ is given as

$$\pi(\sigma) = \exp\left(-J \sum_{x \sim y} \delta(\sigma(x), \sigma(y))\right) / Z,$$

where the sum is taken over all nearest neighbors in G and Z is the normalizing constant or partition function. In the unsaturated setting studied here, spins are identified with particles, not sites, and particles can make spatial moves to unoccupied sites in addition to updating

their spins. We present alignment algorithms for two natural variants: (a) the connected setting, where particles are constrained to be simply connected in the lattice, and (b) the general setting, where particles occupy any distinct lattice sites regardless of connectivity.

Recent work on a closely related *site-diluted* Potts model [40, 9] also allows a non-zero fraction of lattice sites to be unoccupied, but the number of particles is not fixed, so particles can appear and disappear, in addition to making spatial moves. Chayes et al. [9] beautifully demonstrate the presence of ordered (aligned and occupied) and disordered (non-aligned and vacant) phases, along with novel “staggered” phases in this model. However, our constraint fixing the number of particles, which is necessary in SOPS models in programmable matter, makes our system fundamentally different from the site-diluted Potts model akin to the difference between the fixed magnetization Ising model, which has a fixed number of + spins, and the Ising model in the presence of a magnetic field, where the number of + spins can vary. Notably, the coexistence of phases that characterize the aligned and compressed behaviors we are seeking will not occur unless we fix the magnetization (or numbers of particles) as these configurations are exponentially unlikely in the site-diluted model and thus do not inherit any of its properties.

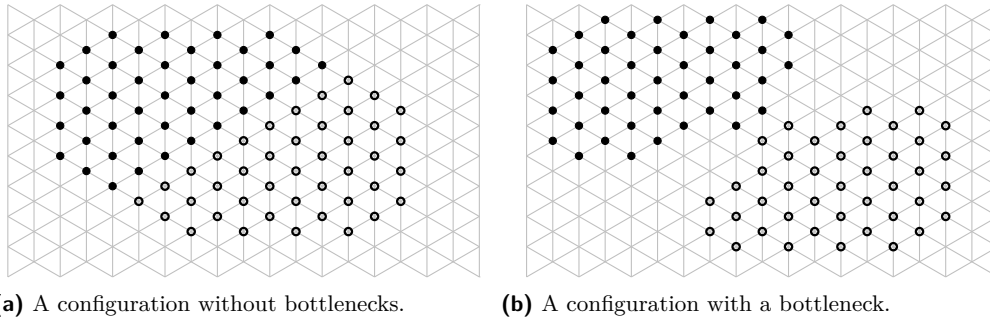
Since particles can make spatial moves, the boundary between the particle occupied sites and the unoccupied sites can assume arbitrary shapes, which makes achieving alignment more challenging than achieving compression. Consider the configurations shown in Figure 1(a),(b), where the particles can be oriented along one of two possible directions ($q = 2$) shown by black and grey circles, with a total of n particles. While the number of unaligned pairs of adjacent particles is $O(\sqrt{n})$ for the configuration in Figure 1(a), it can be as low as $O(1)$ for the configuration shown in Figure 1(b), owing to the bottleneck shaped part of the configuration boundary, making it likely that the regions on either side of it will be aligned along different directions. Hence, achieving alignment requires suppressing the likelihood of such bottlenecks in the boundary of the particle configuration.

While the concept of an interfacial free energy can be used to constrain the shape of the boundary of a dilute system of homogeneous particles i.e., when $q = 1$, as in [26, 17, 2, 30], because particle occupied sites and vacant sites are akin to distinct coexisting phases of the system. However, the same ideas do not readily generalize to the case when $q \geq 2$. Instead, we show build on the notion of compression introduced in [7, 8], and use isoperimetric inequalities to show that for sufficiently compressed configurations, bottlenecks such as the one shown in Figure 1(b) are precluded with high probability.

1.2 Results

We present the first rigorous local distributed algorithms for achieving both low perimeter boundaries and alignment, for any number of orientations $q \geq 2$, in both connected and general settings. Informally, we say a particle system is *aligned* if a significant percentage of the particles have the same orientation.

In the connected SOPS setting, we define an energy function that encourages compression of the entire configuration and also defines a ferromagnetic interaction between particles’ orientations, inspired by the clock and Potts models. These two contributions are controlled by two independent parameters λ and γ . In this setting, we show that given any $\alpha > 1$, for any $\lambda > 1$ and $\gamma > 29.3(q - 1)$ such that $\lambda\gamma > 7^{\alpha/(\alpha-1)}$, the algorithms achieve α -compression with high probability. Furthermore, when γ satisfies additional constraints given in Theorem 11, we show that the compressed configurations are very likely to be aligned. Next, we show that setting λ large and γ small will generate compressed configurations



■ **Figure 1** Configurations with two dominant orientations (black vs. gray circles); large interfaces as in (a) are unlikely for large γ , whereas small interfaces as in (b) are likely for any finite γ .

with an equitable balance of orientations (Theorem 18), while setting λ small will generate configurations that are *expanded*, nearly maximizing their perimeters, allowing the SOPS to explore space, potentially to forage for resources, for example (see Theorem 19).

For both the Potts and clock models in the connected setting, the proofs rely on the *cluster expansion* [23, 14, 20] from statistical physics, introducing a new so-called *polymer model* inspired by the relationship between flows and the Potts model [13]. Informally, the cluster expansion allows us to obtain upper and lower bounds on the so-called “polymer partition function” in terms of the volume and surface contributions, as in [5, 6, 14], to prove that our algorithms achieve compression (or aggregation), with high probability. Moreover, using isoperimetric inequalities, we prove the absence of bottlenecks in sufficiently highly compressed configurations, which is necessary to get the system to globally align. Finally, we use the *bridging techniques* first proposed in [27] and later adapted in [5, 6], to expand the information theoretic arguments in [5, 6] to prove that for sufficiently compressed configurations, our algorithms achieve alignment with high probability. Conversely, we show that our algorithms can achieve *expansion* and/or *non-alignment* (with all directions nearly equitably balanced), with the same algorithm by adjusting only two global parameters.

In the general SOPS setting, with no connectivity constraints, we present an algorithm based on a single parameter coupling both compression and ferromagnetism simultaneously. When this parameter is sufficiently large, we achieve aggregation and alignment, while when it is small we achieve expansion and a balance among the orientations (Theorem 21). We believe these parameters can be independently controlled in the general (disconnected) setting, but the proofs seemingly become significantly more challenging and coupling them into one parameter seems sufficient for most applications in programmable matter and swarm robotics. Because configurations tend to be highly disconnected, proofs in the general setting require additional technology to account for many small clusters that can be distributed throughout the lattice. Here we generalize the bridging techniques to account for more complex contours that form an interconnected network to show that the contour lengths of the bridging system can be made arbitrarily close to their minimum possible length and, as a result, alignment occurs with high probability. We note that our algorithms for alignment in both settings work for all $q \geq 2$; separation (where the sizes of the color classes are fixed) has only been shown for $q = 2$, although the methods should also generalize to more colors [6].

2 Preliminaries

Our model of programmable matter is based on the *amoebot model*, introduced in [11] and described in detail in [10], which has served as the basis for previous stochastic algorithms for SOPS [8, 7, 6, 5]. In the amoebot model, particles occupy the nodes of a graph with each node occupied by at most one particle. When executing a spatial move, a particle expands into an adjacent unoccupied node, temporarily occupying both nodes and then contracts to the new node. Each particle stores whether it is expanded or contracted and can read whether its neighbors are expanded or contracted. No particle has access to global information such as system size or a shared co-ordinate system or compass.

We extend the amoebot model to model heterogeneous particles, where each particle has one of q orientations, akin to the variant introduced in [6, 5]. Each particle, when activated, chooses either a spatial move as in the original amoebot model, or an “orientation move” that updates its direction, each equal probability. The system performs these *atomic actions*, following the *ASync* model of computation from distributed computing [22]. It has been shown in this model that for any concurrent asynchronous execution of atomic actions, there exists a sequential ordering of actions with the same end state provided that all conflicts arising in the concurrent asynchronous execution are resolved. We assume that conflicts due to multiple particles expanding into an unoccupied node are resolved arbitrarily so that only one particle expands into the unoccupied node, allowing us to consider only one particle to be active at any given time.

2.1 The Potts and clock models

In our models, each configuration is an assignment of n particles to distinct vertices of a finite triangular lattice G_Δ of $N > n$ vertices with the toroidal topology. In addition, each particle is also assigned an orientation from $\{0, 1, \dots, q-1\}$. We assume G_Δ to inhabit a $\sqrt{N} \times \sqrt{N}$ square region with periodic boundary conditions. Each vertex (x, y) of G_Δ has six outgoing edges, to the vertices $(x+1, y)$, $(x, y+1)$, $(x+1, y+1)$, $(x-1, y)$, $(x, y-1)$, $(x-1, y-1)$, where addition and subtraction is taken modulo $\sqrt{N}-1$. Moreover, in this setup, the set of particles in our configurations must always be connected and hole-free. Given such a configuration, we define its *boundary* \mathcal{P} to be the minimal closed walk over occupied sites of G_Δ that encloses all of the occupied sites in the configuration. The *perimeter* $p(\sigma)$ of a configuration σ is then defined to be the length of this closed walk.

We consider the following Potts Hamiltonian, on G_Δ , a variant of the site-diluted Potts model [9]:

$$H_{\text{Potts}}(\sigma) = -J \sum_{\langle i, j \rangle} n_i n_j \delta(\theta_i, \theta_j) - \kappa \sum_{\langle i, j \rangle} n_i n_j,$$

where the sum is over all pairs of adjacent sites: $\langle i, j \rangle$ i.e., sites connected by a single lattice edge in G_Δ , $n_i \in \{0, 1\}$ indicates whether site i is occupied or not, θ_i indicates the orientation of the particle on site i , and J, κ are positive constants. We only consider configurations σ in Ω , i.e., where the total number of particles is equal to n , and the particle-occupied sites form a connected, hole-free region.

The probability of a configuration $\pi_{\text{Potts}}(\sigma)$ is given by the Boltzmann distribution:

$$\pi_{\text{Potts}}(\sigma) = e^{-\beta H_{\text{Potts}}(\sigma)} / Z_{\text{Potts}}, \quad \text{where} \quad Z_{\text{Potts}} = \sum_{\sigma' \in \Omega} e^{-\beta H_{\text{Potts}}(\sigma')},$$

where β denotes the inverse temperature. Setting parameters $\lambda = \exp(\beta\kappa)$, and $\gamma = \exp(\beta J)$, the above probability distribution can be expressed as:

$$\pi_{\text{Potts}}(\sigma) = \frac{w_{\text{Potts}}(\sigma)}{Z_{\text{Potts}}}, \quad w_{\text{Potts}}(\sigma) = (\lambda \gamma)^{-p(\sigma)} \gamma^{-h(\sigma)}, \quad Z_{\text{clock}} = \sum_{\sigma' \in \Omega} w_{\text{Potts}}(\sigma'), \quad (1)$$

where $h(\sigma)$ is the number of heterogeneous edges in σ , i.e., edges connecting particles with different orientations, and $p(\sigma)$ is its perimeter, as defined earlier. Here π_{Potts} is the stationary distribution for our Markov chain algorithm based on the ferromagnetic Potts model interactions.

Similarly, we consider the following clock model Hamiltonian on G_{Δ} :

$$H_{\text{clock}}(\sigma) = -J \sum_{\langle i,j \rangle} n_i n_j \cos(2\pi(\theta_i - \theta_j)/q) - \kappa \sum_{\langle i,j \rangle} n_i n_j.$$

The probability of a configuration $\pi_{\text{clock}}(\sigma)$ is given by the Boltzmann distribution as before, and can be expressed in terms of the parameters λ, γ as:

$$\pi_{\text{clock}}(\sigma) = \frac{w_{\text{clock}}(\sigma)}{Z_{\text{clock}}}, \quad w_{\text{clock}}(\sigma) = (\lambda \gamma)^{-p(\sigma)} \prod_{\langle i,j \rangle} \gamma^{-d_{ij}}, \quad Z_{\text{clock}} = \sum_{\sigma' \in \Omega} w_{\text{clock}}(\sigma'), \quad (2)$$

where $\lambda > 0, \gamma > 0$ (as before), $d_{ij} := 1 - \cos(2\pi(\theta_i - \theta_j)/q)$, and the product is over all pairs of adjacent occupied sites. Here π_{clock} will be the stationary distribution for our Markov chain algorithm based on the clock model.

For each of the above models, we will refer to $w(\sigma)$ (w_{Potts} or w_{clock}) as the *weight* of a configuration. The stationary probability distribution π (π_{Potts} or π_{clock}) is thus simply the weight function w normalized by the *partition function* Z (Z_{Potts} or Z_{clock}).

2.2 Cluster expansions and bridging

Our proofs build on several tools from statistical physics and combinatorics, so we begin by introducing two key methods. The cluster expansion is one of the oldest tools in statistical physics [23, 24, 14], and has led to the development of the Pirogov-Sinai theory [31, 32], playing an important role in recent advances in efficient sampling and counting algorithms [16, 19, 3]. The cluster expansion expresses the logarithm of a polymer partition function as a sum over polymer clusters.

Let \mathcal{L} be a finite set of polymers $\{\xi_i\}$, where each polymer ξ_i has weight $w(\xi_i)$. We also define ‘‘compatibility’’ between polymers - each pair of polymers ξ, ξ' is either compatible ($\xi \sim \xi'$) or incompatible ($\xi \approx \xi'$). The polymer partition function is then given by:

$$\Xi = \sum_{\tau \in \Omega^{\mathcal{L}}} \prod_{\xi \in \tau} w(\xi),$$

where $\Omega^{\mathcal{L}}$ is the set of all collections of pairwise compatible polymers in \mathcal{L} . The cluster expansion expresses the logarithm of the polymer partition function in terms of clusters, where a cluster X is an ordered multiset of polymers $\{\xi_1, \dots, \xi_k\}$ such that their incompatibility graph $H(X)$ is connected, where the incompatibility graph is constructed by representing each polymer by a vertex and connecting two vertices if the corresponding polymers are incompatible. The cluster expansion gives:

$$\log \Xi = \sum_{X \in \mathcal{C}} \Psi(X), \quad \text{where } \Psi(X) := \frac{1}{|X|!} \left(\sum_{G \subseteq H_X} (-1)^{|E(G)|} \right) \left(\prod_{\xi \in X} w(\xi) \right),$$

where the sum is taken over connected, spanning subgraphs G and \mathcal{C} is the set of all clusters. A sufficient condition for the convergence of the cluster expansion was given by Kotecký and Preiss [20]. We will prove this condition in Lemma 7 and use the cluster expansion to separate the volume and surface contributions to the partition function, as done in [14, 6].

Bridging is a combinatorial technique used to show that large contours are uncommon, while allowing for the possibility of many small contours corresponding to “defects”. It was first introduced in [27] and later adapted in [6]. We note that a constant fraction of defects will be unavoidable - an example of this is in the Ising model and Potts models, where a constant fraction of the vertices will not follow the majority color even at stationarity. Each configuration corresponds to a set of contours - informally, a bridge system comprises of a set of bridges, which are edges on the dual graph on the lattice that connect contours to the boundary of the lattice. Contours that are connected this way are called bridged contours, while the remaining contours are unbridged.

Bridge systems are defined so that the total length of the bridges is at most a constant fraction of the total length of the bridged contours, which allows us to bound the number of bridge systems with total bridged contour length ℓ by C^ℓ for some constant C . Consequently, a *Peierls argument* can be used to show that the gain in energy (probability weight) by the removal of the bridged contours is greater than the loss in entropy by the removal of these contours. Explicit constructions of bridge systems are shown in [6] and in our proof of alignment for disconnected SOPS (see Appendix A).

3 Compression and Alignment in Connected SOPS

Starting with any simply connected set of particles, we define a local Markov chain aiming to simultaneously compresses the configuration and align all but a small fraction of their orientations. On each iteration, a particle is activated uniformly at random using a Poisson clock. When activated, a particle chooses to attempt a spatial move or a reorientation move with a equal probability. Informally, spatial moves consist of the particle moving to a randomly chosen neighboring site, provided that site is unoccupied and the particle configuration remains simply connected, while a reorientation move allows the particle to change its orientation to point in a new direction. While it is surprising that a property such as connectivity can be determined locally, a set of local moves were defined in Cannon et al. [8] that prevent the configuration from disconnecting or forming holes and yet the chain remains ergodic on the infinite lattice, so all valid configurations can still be reached. This ergodicity result carries over to our setting as the we use a lattice that while finite, is sufficiently large that self-intersections via wraparound are not possible. Using the Metropolis-Hastings algorithm [25], once a move is determined to be valid, it is implemented with probability $\min\{1, \pi(\sigma')/\pi(\sigma)\}$, where π is the desired stationary distribution.

More precisely, consider a spatial move from a location ℓ to an empty adjacent location ℓ' . Let the sets of lattice sites adjacent to the locations ℓ and ℓ' be $N(\ell)$ and $N(\ell')$ respectively. Furthermore, let $N(\ell \cup \ell')$ denote $N(\ell) \cup N(\ell') \setminus \{\ell, \ell'\}$, and $\mathbb{S} := N(\ell) \cap N(\ell')$ denote the set of sites adjacent to both ℓ and ℓ' so that $|\mathbb{S}| \in \{0, 1, 2\}$.

► **Definition 1.** *A move from ℓ to ℓ' is valid if ℓ' is unoccupied, the number of particle-occupied sites in $N(\ell)$ is less than 5, and either of the following two properties are satisfied:*

Property 1: $|\mathbb{S}| \geq 1$ and every particle-occupied site in $N(\ell \cup \ell')$ is connected to a particle-occupied site in \mathbb{S} through $N(\ell \cup \ell')$.

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

Note that in Section 4, we will consider almost the same algorithm in the general SOPS setting where there are no connectivity restrictions, so there all spatial moves from an occupied site to an adjacent unoccupied site are valid.

It is important to note that the ratio between the probabilities $\pi(\sigma')/\pi(\sigma)$ that arises from the Metropolis-Hastings algorithm can be calculated by an activated particle using only local information - the positions and orientations of particles in its immediate neighborhood, as well as those in the neighborhood of the destination site if the particle is moving. Specifically, changes in perimeter in connected SOPS can be computed locally as shown in [8, 7].

We now proceed to show that when s , λ and γ are sufficiently large, the alignment algorithm will cause the system to compress to form a low-perimeter configurations with high probability. Moreover, in both the Potts and clock model settings, in any configuration with sufficiently low-perimeter, one of the q orientations will dominate with high probability.

We note that we did not attempt to give rigorous bounds on the rates of convergence for our Markov chains. We expect that convergence will be fast when the parameters λ and γ are small and the system evolves to a disordered (gaseous) state, but the connectivity constraint makes proving this challenging. In contrast, we expect convergence to equilibrium will be slow in the ordered (solid) state when λ is large, but we conjecture that desirable compressed and aligned states will be reached quickly, long before the system is very close to stationarity.

3.1 Compression in Connected SOPS

We denote the set of possible configurations in this paradigm by Ω . Recall that N represents the number of sites of the lattice G_Δ . To ensure that the proof of ergodicity from [7] carries over to our setting, we use a sufficiently large value of N , namely $N \geq (n+1)^2$, although we expect the results to hold for smaller N .


► **Definition 2 (Compression).** *A simply connected configuration σ of n particles on a lattice is said to be α -compressed if its perimeter is at most $\alpha \cdot p_{\min}(n)$, where $p_{\min}(n)$ is the minimum possible perimeter of a configuration of n particles.*

The main result of this section is the following theorem.

► **Theorem 3.** *Given any $\alpha > 1$, if constants $\lambda > 1$ and $\gamma > 29.3(q-1)$ satisfy $\lambda\gamma > 7^{\alpha/(\alpha-1)}$ and n is sufficiently large, then the probability a configuration drawn from the stationary distribution π_{Potts} is not α -compressed is exponentially small.*

Let \mathcal{P} denote the boundary of some configuration σ in our configuration space Ω . As σ is connected, hole-free, and contains a finite (n) number of particles, \mathcal{P} is a single closed walk on G_Δ and the perimeter of the configuration, $p(\sigma)$, is equal to $|\mathcal{P}|$, the total length of walk \mathcal{P} . If we restrict our particle configurations to be connected and hole-free, there is a one-to-one correspondence between the possible sets of occupied sites and the possible boundaries \mathcal{P} . Let $\Omega_{\mathcal{P}}$ denote the set of configurations in Ω with boundary \mathcal{P} , and let $\Lambda_{\mathcal{P}} \subseteq G_\Delta$ be the induced subgraph of the triangular lattice G_Δ by the particle-occupied vertices for any configuration in $\Omega_{\mathcal{P}}$. A configuration in $\Omega_{\mathcal{P}}$ thus corresponds to a mapping of the vertices of $\Lambda_{\mathcal{P}}$ to the orientations $\{0, \dots, q-1\}$.

We consider the subset of configurations $\Omega_{\mathcal{P}}^0 \subseteq \Omega_{\mathcal{P}}$ where all particles on the boundary \mathcal{P} have the same color 0. We will later analyze the weight of configurations in $\Omega_{\mathcal{P}}^0$ using a polymer model and the cluster expansion. We would first like to obtain an upper bound on $w(\Omega_{\mathcal{P}})$, the total weight of configurations in $\Omega_{\mathcal{P}}$, in terms of $w(\Omega_{\mathcal{P}}^0)$, the total weight of configurations in $\Omega_{\mathcal{P}}^0$.



■ **Figure 2** Particle configuration in $\Omega_{\mathcal{P}}^0$, and its corresponding polymer configuration in $\Omega_{\mathcal{P}}^{\mathcal{L}}$ (with two polymers).

► **Lemma 4.** For $\gamma > 3q$, we have

$$w(\Omega_{\mathcal{P}}) < w(\Omega_{\mathcal{P}}^0) \cdot q 2^{|\mathcal{P}|} \frac{\gamma}{\gamma - 3q}.$$

The proof is a generalized version of that in [6], by defining maps from $\Omega_{\mathcal{P}} \rightarrow \Omega_{\mathcal{P}}^0$ such that all vertices on boundary \mathcal{P} are of orientation 0. We will use the cluster expansion to analyze the total weight $w(\Omega_{\mathcal{P}}) := \sum_{\sigma \in \Omega_{\mathcal{P}}} w(\sigma)$ of the configurations in $\Omega_{\mathcal{P}}$. Since the cluster expansion can only be applied to polymer partition functions, we begin by representing the configurations of $\Omega_{\mathcal{P}}$ with a polymer model.

The Polymer Model. We say two edges of G_{Δ} are adjacent if they share a common vertex. A polymer ξ in \mathcal{L} is defined to be a labeling $\xi : E(G_{\Delta}) \rightarrow \{0, 1, \dots, q - 1\}$ of the edges of G_{Δ} such that the set $E(\xi)$, defined to be the edges of G_{Δ} with a non-zero label in ξ , is non-empty and connected under the above notion of adjacency. The labeling must also be *consistent*, as defined below.

► **Definition 5 (Consistent Labeling).** We fix a canonical direction for each edge in G_{Δ} . This direction can be arbitrarily defined, so for simplicity we say that the edge is oriented toward the vertex with the larger x , followed by y coordinate, where the coordinate axes are oriented such that the x coordinate increases from left to right and the y coordinate increases from top to bottom.

We define labels $\xi : E(G_{\Delta}) \rightarrow \{0, 1, \dots, q - 1\}$. These edge labels represent “flows” in our defined canonical direction, modulo q . In other words, when summing up the total flow along a walk on G_{Δ} , for each edge e on the walk, we add the label $\xi(e)$ to the sum if the walk is in the canonical direction of the edge, and $q - \xi(e)$ if the walk is in the opposite direction. We call an assignment of labels consistent if every closed walk on G_{Δ} has a total flow summing to 0 modulo q .

Consider a fixed boundary \mathcal{P} as defined above, corresponding to some configuration in Ω . For a polymer ξ , denote by $V(\xi)$ the set of vertices incident to an edge with a non-zero label in ξ . We say a polymer ξ is within \mathcal{P} if $V(\xi) \subseteq \Lambda_{\mathcal{P}}$. As described earlier, the set $\Omega_{\mathcal{P}}^{\mathcal{L}}$ of polymer configurations corresponding to \mathcal{P} is the set of all subsets of \mathcal{L} of pairwise compatible polymers within \mathcal{P} . The weight $w(\tau)$ of a configuration $\tau \in \Omega_{\mathcal{P}}^{\mathcal{L}}$ is the product of the weights of its constituent polymers.

Two polymers ξ_1, ξ_2 are incompatible if there are edges $e_1 \in E(\xi_1)$ and $e_2 \in E(\xi_2)$ such that e_1 and e_2 are adjacent. The weight of a polymer ξ is defined as $w(\xi) := \gamma^{-|E(\xi)|}$, in the Potts model, and $w(\xi) := \prod_{e \in E(\xi)} \gamma^{\cos(\frac{2\pi}{q}\xi(e)) - 1}$ in the clock model.

► **Lemma 6.** There is a bijection ϕ between $\Omega_{\mathcal{P}}^0$ and $\Omega_{\mathcal{P}}^{\mathcal{L}}$ with the property that for any $\sigma \in \Omega_{\mathcal{P}}^0$, we have $w(\sigma) = (\lambda\gamma)^{-P(\sigma)} w(\phi(\sigma))$.

14:10 Local Stochastic Algorithms for Alignment in Self-Organizing Particle Systems

The map ϕ simply encodes the orientations of particles in a configuration $\sigma \in \Omega_{\mathcal{P}}^0$ as differences between orientations on the edges of G_{Δ} . This is illustrated in Figure 2. The full version of the paper gives a full description of this mapping and a proof that it is indeed a bijection. From Lemma 6, we have

$$w(\Omega_{\mathcal{P}}^0) = \sum_{\sigma \in \Omega_{\mathcal{P}}^0} (\lambda\gamma)^{-|\mathcal{P}|} w(\phi(\sigma)) = \sum_{\tau \in \Omega_{\mathcal{P}}^{\mathcal{L}}} (\lambda\gamma)^{-|\mathcal{P}|} w(\tau) = (\lambda\gamma)^{-|\mathcal{P}|} \Xi_{\mathcal{P}},$$

where $\Xi_{\mathcal{P}}$ is the partition function for the set of polymer configurations $\Omega_{\mathcal{P}}^{\mathcal{L}}$:

$$\Xi_{\mathcal{P}} := \sum_{\tau \in \Omega_{\mathcal{P}}^{\mathcal{L}}} w(\tau) = \sum_{\tau \in \Omega_{\mathcal{P}}^{\mathcal{L}}} \prod_{\xi \in \tau} w(\xi).$$

The Potts Model. From now, our analysis will be specific to the Potts model. The clock model will be discussed in Section 3.1. The following Lemmas and proofs are slight variations of those used in [6].

► **Lemma 7.** *For any polymer $\xi \in \mathcal{L}$, whenever $\gamma > 29.3(q-1)$, we have for $c = 0.0001$,*

$$\sum_{\substack{\xi' \in \mathcal{L} \\ \xi' \approx \xi}} w(\xi') \exp(c|V(\xi')|) \leq c|V(\xi)|,$$

where $V(\xi')$ denotes the set of vertices in the polymer ξ' , and $|V(\xi')|$ denotes the number of vertices in ξ' .

The proof is on the lines of that in [6]. The key part of this proof is the use of an upper bound $\nu(m, q) \leq (6e(q-1))^m/2$ from [4], where $\nu(m, q)$ represents the number of polymers with m edges containing some fixed vertex $v \in V(G_{\Delta})$.

Lemma 7 has an important consequence in addition to guaranteeing the convergence of the cluster expansion, as stated in the original paper of Kotecký and Preiss [20], and rephrased in [19]. Consider the function $\Psi(X)$ defined earlier for any cluster X . An additional consequence [20, 19] of Lemma 7 is that $\Psi(X)$ will satisfy the following inequality

$$\sum_{\substack{X \in \mathcal{X} \\ X \approx \xi}} |\Psi(X)| \leq c|V(\xi)|. \quad (3)$$

for any polymer ξ , where \mathcal{X} is the set of all clusters of polymers, and a cluster $X \approx \xi$ if there exists a polymer $\xi' \in X$ such that $\xi' \approx \xi$. The support of a cluster X is denoted by \bar{X} and is given by $\bar{X} = \bigcup_{\xi \in X} V(\xi)$.

Consider an arbitrary vertex $v \in G_{\Delta}$, and let ξ_v be the smallest polymer consisting of six edges of equal weight attached to v . From Equation (3), we have:

$$\sum_{\substack{X \in \mathcal{X} \\ X \approx \xi_v}} |\Psi(X)| \leq c|V(\xi_v)| = 7c \Rightarrow \sum_{\substack{X \in \mathcal{X} \\ v \in \bar{X}}} |\Psi(X)| \leq \sum_{\substack{X \in \mathcal{X} \\ X \approx \xi_v}} |\Psi(X)| \leq 7c. \quad (4)$$

► **Lemma 8.** *If for any polymer $\xi \in \mathcal{L}$, there exists a constant c such that*

$$\sum_{\substack{\xi' \in \mathcal{L} \\ \xi' \approx \xi}} w(\xi') \exp(c|V(\xi')|) \leq c|V(\xi)|,$$

then for any connected region $\Lambda_{\mathcal{P}}$ with boundary \mathcal{P} , the partition function $\Xi_{\mathcal{P}}$ satisfies

$$\psi|\Lambda_{\mathcal{P}}| - 7c|\partial\Lambda| \leq \ln \Xi_{\mathcal{P}} \leq \psi|\Lambda_{\mathcal{P}}| + 7c|\partial\Lambda|.$$

The proof follows on the lines of the proof of a similar Lemma in [6], and section 5.7.1 of [14]. Using Lemma 8, and noting that $|\partial\Lambda_{\mathcal{P}}| \leq p(\sigma) \forall \sigma \in \Omega_{\mathcal{P}}$ and $|\Lambda_{\mathcal{P}}| = n$, we get:

$$n\psi - 7cp(\sigma) \leq \ln \Xi_{\mathcal{P}} \leq n\psi + 7cp(\sigma) \tag{5}$$

Note that the partition function Z_{Potts} is greater than the contribution from particle configurations in $\Omega_{\mathcal{P}}^0$ where the length of the boundary is the smallest attainable perimeter $|\mathcal{P}| = p_{\min}$:

$$Z_{\text{Potts}} \geq w(\Omega_{\mathcal{P}}^0) = (\lambda\gamma)^{-p_{\min}} \Xi_{\mathcal{P}} \geq (\lambda\gamma)^{-p_{\min}} e^{n\psi - 7cp_{\min}}. \tag{6}$$

Given $\alpha > 1$, let S_{α} be all configurations that are not α -compressed. We will prove that the probability of the set S_{α} in the stationary distribution is exponentially small for sufficiently large λ, γ :

► **Lemma 9.** *Given any $\alpha > 1$, when constants $\lambda > 1, c = 0.0001$, and $\gamma > 29.3(q - 1)$ satisfy*

$$\lambda\gamma > (4 + 2\sqrt{2})^{\frac{\alpha}{\alpha-1}} (e^{7c})^{\frac{\alpha+1}{\alpha-1}} \tag{7}$$

and n is sufficiently large, then the probability that a configuration drawn from the stationary distribution π_{Potts} is not α -compressed is exponentially small, $\pi_{\text{Potts}}(S_{\alpha}) < \zeta^{\sqrt{n}}$.

Note that Equation (7) is satisfied if $\lambda\gamma > 7^{\alpha/(\alpha-1)}$, proving Theorem 3. The proof of the Lemma requires using Lemma 4, Lemma 8 and Equation (6), and an upper bound on the number of self-avoiding walks of a given length on the triangular lattice from [12, 8].

The Clock Model. The proof of compression for the clock-model-inspired algorithm follows along the same lines as the proof for the Potts-model-inspired algorithm. The set of allowed particle configurations is the same as before, so the set of configurations in $\Omega_{\mathcal{P}}^0$ is in a one-to-one correspondence with compatible collections of polymers with the same polymer model as above, albeit with the weight of a polymers redefined as $w_{\text{clock}}(\xi) = \prod_{e \in \xi} \gamma^{-d_e}$, where $d_e = 1 - \cos(2\pi\ell(e)/q)$, and $\ell(e) \in \{1, 2, \dots, q - 1\}$ is the label associated with an edge $e \in \xi$. This changes the prefactor in Lemma 4, replacing γ with $\gamma^{-\cos(2\pi/q)}$, and requiring $\gamma^{-\cos(2\pi/q)} > 3q$. The polymer partition function becomes

$$\Xi_{\mathcal{P}} = \sum_{\substack{\mathcal{L}' \subseteq \mathcal{L}_{\mathcal{P}} \\ \text{compatible}}} \prod_{\xi \in \mathcal{L}'} w_{\text{clock}}(\xi).$$

Since the maximum weight of an edge in a polymer is now $\gamma^{-(1-\cos(2\pi/q))}$, instead of γ^{-1} , the condition for Lemma 7 to hold becomes $\gamma^{1-\cos(2\pi/q)} > 29.3(q - 1)$. Lemmas 8 and Theorem 3 follow without modification except for the modified condition: $\gamma^{1-\cos(2\pi/q)} > 29.3(q - 1)$ in Theorem 3.

3.2 Alignment in Compressed Configurations

► **Definition 10 (Alignment).** *We say a configuration of n particles with q orientations is δ -aligned if there exists an orientation $\theta \in \{0, 1, \dots, q - 1\}$, such that the number of particles of orientation θ is at least $(1 - \delta)n$.*

Our main result is the following theorem:

► **Theorem 11.** Denote by $\pi_{\text{Potts}, \mathcal{P}}$ the stationary distribution π_{Potts} conditioned on the boundary of the configuration being \mathcal{P} . For any η where $1/2 < \eta < 1$, there exists a constant $\alpha^* = \alpha^*(\eta, q) > 1$, such that for all α where $1 < \alpha < \alpha^*$, there exists a sufficiently large $\gamma^* = \gamma^*(\eta, q, \alpha, \alpha^*)$ where as long as $\gamma > \gamma^*$ and \mathcal{P} is α -compressed, the probability that a configuration drawn from $\pi_{\text{Potts}, \mathcal{P}}$ is not $(1 - \eta)$ -aligned is exponentially small.

In particular, possible values of α^* and γ^* are:

$$\alpha^*(\eta, q) = \min \left\{ \sqrt{\eta} + \sqrt{1 - \eta}, \sqrt{q^{-1}} + \sqrt{1 - q^{-1}} \right\}$$

$$\gamma^*(\eta, q, \alpha, \alpha^*) = \left(3^{\frac{2\alpha}{\alpha^* - \alpha}} \cdot 4^{\frac{3}{4} + \frac{\alpha^* - 1}{2\delta^*(\eta, q)(\alpha^* - \alpha)}} \right)^{q-1} \text{ where } \delta^*(\eta, q) := \min\{1 - \eta, q^{-1}\}.$$

For any particle configuration, let $2\pi\theta_p/q$, be the *most popular orientation*, or the orientation possessed by the greatest number of particles, where $\theta_p \in \{0, 1, \dots, (q - 1)\}$, and let ρ_p be the fraction of particles with orientation θ_p . Note that $1/q \leq \rho_p \leq 1$, and $\rho_p \geq \eta$ for a $(1 - \eta)$ -aligned configuration.

The dual lattice, G_\square , to the triangular lattice G_Δ is obtained by creating a dual vertex in the center of each triangle in G_Δ , and joining these dual vertices with edges if their corresponding triangular faces share an edge. Each edge e_Δ of G_Δ corresponds with the edge e_\square of G_\square that crosses it. This corresponding edge e_\square separates the two endpoints of e_Δ in G_Δ . A *contour* refers to a self-avoiding walk on the edges of the dual lattice G_\square . The length of a contour refers to the number of edges in the contour.

In this setting, we distinguish between the *boundary contour* and the *internal boundary contour* of a region $R \subseteq V(\Lambda_{\mathcal{P}})$. The boundary contour refers to the set of edges on the dual lattice G_\square corresponding to edges between sites in R and sites not in R , while the internal boundary contour includes edges only from $E(\Lambda_{\mathcal{P}})$ rather than all of $E(G_\square)$. We make use of the following geometric result, which we show in the full version of the paper:

► **Lemma 12.** For a connected hole-free α -compressed configuration with n particles, a particle-occupied region R containing κn particles has an internal boundary contour $bd_{\text{int}}(R)$ of length at least $\nu\sqrt{n}(\sqrt{\kappa} + \sqrt{1 - \kappa} - \alpha)$ for any $\nu < 2\sqrt{3}$ for sufficiently large n .

For the rest of this section, we assign particles the color c_1 if they are of orientation θ_p , and the color c_2 otherwise. This lets us directly apply the bridging construction from [6].

► **Lemma 13** ([6], Lemma 7.3). Fix $\delta \in (0, 1/2)$. For each particle configuration $\sigma \in \Omega_{\mathcal{P}}$, there exists a function $\mathcal{R}_\delta : \Omega_{\mathcal{P}} \rightarrow 2^{\Omega_{\mathcal{P}}}$ giving a region $\mathcal{R}_\delta(\sigma)$ such that all particles on the boundary of $\mathcal{R}_\delta(\sigma)$ have the color c_1 , all particles on the boundary of its complement $\bar{\mathcal{R}}_\delta(\sigma)$ have the color c_2 , $\mathcal{R}_\delta(\sigma)$ contains at most δ fraction of particles with the color c_2 , and $\bar{\mathcal{R}}_\delta(\sigma)$ contains at most δ fraction of particles with the color c_1 .

We use the bridging construction from [6] to define the region $\mathcal{R}_\delta(\sigma)$ in Lemma 13.

► **Lemma 14.** For any particle configuration $\sigma \in \Omega_{\mathcal{P}}$ with total number of particles n and ρ_p fraction of particles of color c_1 , given any $\delta > 0$, the region $\mathcal{R}_\delta(\sigma)$ defined in Lemma 13 is such that the number of particles in $\mathcal{R}_\delta(\sigma)$, $n_{\mathcal{R}_\delta}$ satisfies: $(\rho_p - \delta)n/(1 - \delta) \leq n_{\mathcal{R}_\delta} \leq (\rho_p n)/(1 - \delta)$.

The proof of Lemma 14 follows from noting that the particles in $\mathcal{R}_\delta(\sigma)$ and $\bar{\mathcal{R}}_\delta(\sigma)$ are predominantly of the colors c_1 and c_2 respectively, with an error fraction bounded by δ , and enforcing that the total number of particles with the color c_1 is $\rho_p n$.

► **Lemma 15.** *For a connected hole-free α -compressed configuration $\sigma \in \Omega_{\mathcal{P}}$ that is not $(1 - \eta)$ -aligned for some $\eta < 1$, given any δ where $0 < \delta < \min\{q^{-1}, 1 - \eta\}$, the internal boundary contour length $|bd_{\text{int}}(\mathcal{R}_\delta)|$ of the region $\mathcal{R}_\delta(\sigma)$ defined in Lemma 13 obeys the lower bound $|bd_{\text{int}}(\mathcal{R}_\delta)| \geq \nu\sqrt{n}(\alpha_c(\delta, \eta, q) - \alpha)$ for any $\nu < 2\sqrt{3}$ and n sufficiently large, where*

$$\alpha_c(\delta, \eta, q) := \min \left\{ \sqrt{\frac{q^{-1} - \delta}{1 - \delta}} + \sqrt{\frac{1 - q^{-1}}{1 - \delta}}, \sqrt{\frac{\eta}{1 - \delta}} + \sqrt{\frac{1 - (\eta + \delta)}{1 - \delta}} \right\}.$$

Lemma 15 is a direct consequence of Lemmas 14 and 12. Given an α -compressed boundary \mathcal{P} , let $S_{\mathcal{P}}^\eta \subseteq \Omega_{\mathcal{P}}$ be the set of α -compressed configurations with boundary \mathcal{P} that are not $(1 - \eta)$ -aligned for some $\eta < 1$. For each configuration $\sigma \in S_{\mathcal{P}}^\eta$, let $\bar{\mathcal{R}}_\delta(\sigma)$ be the complement of the region $\mathcal{R}_\delta(\sigma)$ defined in Lemma 13.

Let $\mathcal{P}_{\bar{\mathcal{R}}_\delta}^{\text{int}}$ denote the walk on the edges of G_Δ , each of whose endpoints is a particle in $\bar{\mathcal{R}}_\delta(\sigma)$ that is connected by an edge in G_Δ to a particle in $\mathcal{R}_\delta(\sigma)$. Let $\Theta_{\bar{\mathcal{R}}_\delta}^{\text{int}}$ denote the set of orientations of particles that are incident to an edge in $\mathcal{P}_{\bar{\mathcal{R}}_\delta}^{\text{int}}$, where the orientation of a particle appears as many times as the number of edges connecting that particle to a particle in $\mathcal{R}_\delta(\sigma)$. Note that $|\Theta_{\bar{\mathcal{R}}_\delta}^{\text{int}}| = |bd_{\text{int}}(\bar{\mathcal{R}}_\delta)|$. Let the orientation which appears the most number of times in the set $\Theta_{\bar{\mathcal{R}}_\delta}^{\text{int}}$ be $2\pi\bar{\theta}_p/q$, where $\bar{\theta}_p \in \{0, 1, \dots, q - 1\}$. We consider a map $f_\eta : S_{\mathcal{P}}^\eta \rightarrow \Omega_{\mathcal{P}}$ which applies a cyclic shift to the orientations of all particles in $\bar{\mathcal{R}}_\delta(\sigma)$, so that under f_η , a particle orientation θ is mapped to $(\theta + (\theta_p - \bar{\theta}_p)) \pmod{q}$. Note that this transformation maps the orientation $\bar{\theta}_p$ to θ_p .

► **Lemma 16** ([6]). *For a configuration $\tau \in \text{Im}(f_\eta(S_{\mathcal{P}}^\eta))$, the number of preimages $\sigma \in S_{\mathcal{P}}^\eta$ for which $|bd_{\text{int}}(\mathcal{R}_\delta(\sigma))| = \ell$, where $\mathcal{R}_\delta(\sigma)$ is defined in Lemma 14, is at most $q^3 |P| 4^{\frac{1+3\delta}{4\delta}\ell}$.*

The proof follows from Lemma 7.6 in [6] and by noting that once the internal boundary contour of $\mathcal{R}_\delta(\sigma)$ is known, one of q cyclic shifts in $\bar{\mathcal{R}}_\delta(\sigma)$ recovers σ , given τ .

In this section so far, our results were valid for both the Potts and the clock models. We now consider specifically the case of the Potts model with stationary distribution π_{Potts} . Using the definition of f_η , we find the following.

► **Lemma 17.** *For a configuration $\sigma \in S_{\mathcal{P}}^\eta$, let region $\mathcal{R}_\delta(\sigma)$ be defined as in Lemma 13 with $|bd_{\text{int}}| = \ell$. For the new configuration $f_\eta(\sigma)$ under the map f_η , the ratio $w(\sigma)/w(f(\sigma))$ is at most $(1/\gamma)^{\ell/(q-1)}$.*

The proof of Theorem 11 follows from an information theoretic argument similar to that in [6], by showing that the minimum gain in the weight of a configuration under the map f_η outweighs the maximum number of preimages of the map, and using Lemma 15 to get a lower bound on the gain under f_η . A key component is ensuring that it is possible to choose the parameter $0 < \delta < q^{-1}$, so that the conditions on α and γ described in the theorem statement can be simultaneously satisfied.

The Clock Model. Lemma 17 and Theorem 11 hold for the clock model with stationary distribution π_{clock} , with γ replaced by $\gamma^{1 - \cos(2\pi/q)}$ in both. The proofs follow on similar lines as for the Potts model.

3.3 Non-Alignment and Expansion in Connected SOPS

An interesting artifact of the alignment algorithm is that when λ, γ are small, the opposite properties are achieved, namely nonalignment and expansion. We outline the main results.

Non-alignment in compressed configurations. For $\epsilon > 0$, we say a configuration is ϵ -non-aligned if the fraction of particles of each orientation is within an ϵ -neighborhood of q^{-1} . Let $S_{\mathcal{P}}^{\epsilon}$ denote the set of configurations which have perimeter \mathcal{P} and are not ϵ -non-aligned, and let S^{ϵ} be the set of configurations that are not ϵ -non-aligned. Our main result is as follows:

► **Theorem 18.** *When $\gamma > 0$ satisfies:*

$$\gamma^3 < \left(1 - \epsilon \frac{q}{q-1}\right)^{\frac{q-1}{q} - \epsilon} (1 + \epsilon q)^{\frac{1}{q} + \epsilon} = 1 + \frac{\epsilon^2 q^2}{q-1} + O(\epsilon^3),$$

the probability that a configuration sampled from the stationary distribution of the Markov chain algorithm π_{Potts} is not ϵ -non-aligned is exponentially small, for sufficiently large n .

The proof follows from Stirling's approximation [35] for the number of configurations that are not ϵ -non-aligned, and using rough lower and upper bounds on the weight of configurations in $\Omega_{\mathcal{P}}$. The result also holds for the clock model with γ replaced with γ^2 .

Expansion in Connected SOPS. We define the notion of expansion, on the lines of [8], as follows. We say a configuration σ is β -expanded when its perimeter $p(\sigma)$ is greater than βp_{\max} , where $0 < \beta < 1$. Consider the set of configurations S_{β} that are not β -expanded. Our main result is:

► **Theorem 19.** *For constants $\lambda, \gamma > 0, c_1 = 2.17, c_2 = 2 + \sqrt{2}$ such that $\lambda \gamma^{5/2} < c_1$, and for any β such that:*

$$0 < \beta < \frac{\log c_1 - \log \lambda - \frac{5}{2} \log \gamma}{\log c_2 - \log \lambda - \log \gamma},$$

the probability that a configuration drawn from the stationary distribution π is not β -expanded is exponentially small.

We can get rough upper and lower bounds for the weight of configurations in $\Omega_{\mathcal{P}}$ by estimating the number of ways of getting a fixed perimeter using the bounds in [12, 8].

The same theorem holds for the clock model, with $\gamma^{5/2}$ replaced by γ^4 in the theorem statement, and the proof follows on similar lines.

4 Aggregation and Alignment in General SOPS

In general SOPS, occupying any selection of n out of the N possible sites of G_{Δ} is a valid configuration. Hence, we apply the same Metropolis-Hastings Markov chain as the connected SOPS model, with the exception that any move into an unoccupied location is considered valid regardless of connectivity effects. In this disconnected setting, particles exist on a lattice region with toroidal boundary conditions. We assume the particles occupy a constant fraction ρ of the lattice. Specifically, we define a $\rho \in (0, \frac{1}{3})$ so that $n = \rho N$. The set of possible configurations is denoted $\tilde{\Omega}^{\rho N}$.

Similar to before, *boundary contour* $bd(R)$ of a region $R \subseteq V(G_{\Delta})$ refers to the set of dual edges on G_{\square} corresponding to edges between sites in R and $V(G_{\Delta}) \setminus R$. The *boundary length* of R is $|bd(R)|$. Let $bd_{\min}(k)$ denote the minimum boundary length of a region of k sites in $V(G_{\Delta})$. We restrict ρ to be less than $\frac{1}{3}$ as cases with so many particles (filled sites) that minimum boundary length configurations wrap around the torus G_{Δ} is not instructive for our purposes (a precise explanation for this restriction is in the full version of the paper).

We show that in this general SOPS model, both alignment and aggregation can be achieved with high probability using only local movements. Alignment is defined in Section 3.2, and aggregation is defined as follows:

► **Definition 20** (Aggregation). *For $\alpha > 1$, $\delta > 0$ we say a configuration of n particles is α, δ -aggregated if there exists a region \mathcal{R} such that*

1. *The number of empty sites within \mathcal{R} is at most $\delta|\mathcal{R}|$.*
2. *The number of particles outside of \mathcal{R} is at most $\delta(N - |\mathcal{R}|)$*
3. *The boundary length of \mathcal{R} is at most $\alpha \cdot bd_{\min}(n)$.*

Note that changes in the perimeter of the configuration cannot be locally computed if the set of particles is disconnected. So instead, we make use of the boundary contour length to define our Hamiltonian. More precisely, we consider the following Potts Hamiltonian, another variant of the site-diluted Potts Hamiltonian [9], on G_Δ :

$$\tilde{H}_{\text{Potts}}(\sigma) = -J \sum_{\langle i, j \rangle} [n_i n_j (\delta_{\theta_i, \theta_j} - 1) + (n_i(n_j - 1) + n_j(n_i - 1))] ,$$

where the sum is over all pairs of adjacent sites: $\langle i, j \rangle$ i.e., sites connected by a single lattice edge in G_Δ , $n_i \in \{0, 1\}$ indicates whether site i is occupied or not, θ_i indicates the orientation of the particle on site i , and J is a positive constant. We only consider configurations σ in $\tilde{\Omega}^{\rho N}$ i.e., where the total number of particles is equal to n .

The probability of a configuration $\tilde{\pi}_{\text{Potts}}(\sigma)$ is given by the Boltzmann distribution which can be expressed in terms of the parameter $\lambda = \exp(\beta J)$ as:

$$\tilde{\pi}_{\text{Potts}}(\sigma) = \frac{\tilde{w}_{\text{Potts}}(\sigma)}{\tilde{Z}_{\text{Potts}}}, \quad \tilde{w}_{\text{Potts}}(\sigma) = \lambda^{-a(\sigma) - h(\sigma)}, \quad \tilde{Z}_{\text{Potts}} = \sum_{\sigma' \in \tilde{\Omega}^{\rho N}} \tilde{w}_{\text{Potts}}(\sigma') \quad (8)$$

where $\lambda > 0$, $h(\sigma)$ is the number of heterogeneous edges in the configuration σ , and $a(\sigma)$ is the number of edges between occupied and unoccupied sites in G_Δ .

We prove the following theorem that establishes aggregation and alignment for appropriate settings of the parameters.

► **Theorem 21.** *Fix $\rho < \frac{1}{3}$ and assume that there will always be exactly ρN filled sites on the lattice. For any $\delta > 0$ and $\alpha > 1$, there exists a $\lambda_0 = \lambda_0(q, \rho, \alpha, \delta)$ such that for all $\lambda > \lambda_0$, with probability $1 - \zeta^{\sqrt{N}}$ for some constant $\zeta = \zeta(q, \rho, \alpha, \delta, \lambda) < 1$, there exists a region $\mathcal{R} \subseteq V(G_\Delta)$, where*

1. *There is an orientation $\theta \in \{0, 1, \dots, q-1\}$ where the number of filled sites with orientation θ in \mathcal{R} is at least $(1 - \delta)|\mathcal{R}|$.*
2. *The number of filled sites not in \mathcal{R} is at most $\delta(N - |\mathcal{R}|)$*
3. *The boundary length of \mathcal{R} is at most $\alpha \cdot bd_{\min}(\rho N)$.*

Due to space limitations, we relegate the main details of the proofs to Appendix A.

References

- 1 Logan E. Beaver and Andreas A. Malikopoulos. An overview on optimal flocking. *Annual Reviews in Control*, 51:88–99, January 2021. doi:10.1016/j.arcontrol.2021.03.004.
- 2 T. Bodineau, D. Ioffe, and Y. Velenik. Rigorous Probabilistic Analysis of Equilibrium Crystal Shapes. *Journal of Mathematical Physics*, 41(3):1033–1098, March 2000. arXiv:math/9911106. doi:10.1063/1.533180.

- 3 Christian Borgs, Jennifer Chayes, Tyler Helmuth, Will Perkins, and Prasad Tetali. Efficient sampling and counting algorithms for the Potts model on \mathbb{Z}^d at all temperatures. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 738–751, New York, NY, USA, June 2020. ACM. doi:10.1145/3357713.3384271.
- 4 Christian Borgs, Jennifer Chayes, Jeff Kahn, and László Lovász. Left and right convergence of graphs with bounded degree. *arXiv*, January 2010. arXiv:1002.0115.
- 5 Sarah Cannon, Joshua J. Daymude, Cem Gökmen, Dana Randall, and Andréa W. Richa. A Local Stochastic Algorithm for Separation in Heterogeneous Self-Organizing Particle Systems. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*, volume 145 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 54:1–54:22, Dagstuhl, Germany, 2019. ISSN: 1868-8969. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.54.
- 6 Sarah Cannon, Joshua J. Daymude, Cem Gokmen, Dana Randall, and Andréa W. Richa. A Local Stochastic Algorithm for Separation in Heterogeneous Self-Organizing Particle Systems. *arXiv*, June 2019. URL: <http://arxiv.org/abs/1805.04599>, arXiv:1805.04599.
- 7 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, New York, NY, USA, July 2016. Association for Computing Machinery. doi:10.1145/2933057.2933107.
- 8 Sarah Cannon, Joshua J. Daymude, Dana Randall, and Andréa W. Richa. A Markov Chain Algorithm for Compression in Self-Organizing Particle Systems. *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing - PODC '16*, pages 279–288, 2016. arXiv:1603.07991.
- 9 L. Chayes, R. Kotecký, and S. B. Shlosman. Aggregation and intermediate phases in dilute spin systems. *Commun.Math. Phys.*, 171(1):203–232, July 1995. doi:10.1007/BF02103776.
- 10 Joshua J. Daymude, Kristian Hinnenthal, Andréa W. Richa, and Christian Scheideler. Computing by Programmable Particles. In Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro, editors, *Distributed Computing by Mobile Entities: Current Research in Moving and Computing*, Lecture Notes in Computer Science, pages 615–681. Springer International Publishing, Cham, 2019. doi:10.1007/978-3-030-11072-7_22.
- 11 Zahra Derakhshandeh, Shlomi Dolev, Robert Gmyr, Andréa W. Richa, Christian Scheideler, and Thim Strothmann. 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, June 2014. Association for Computing Machinery. doi:10.1145/2612669.2612712.
- 12 Hugo Duminil-Copin and Stanislav Smirnov. The connective constant of the honeycomb lattice equals $\sqrt{2 + \sqrt{2}}$. *Annals of Mathematics*, 175(3):1653–1665, 2012. Publisher: Annals of Mathematics. URL: <https://www.jstor.org/stable/23234646>.
- 13 J. W. Essam and C. Tsallis. The Potts model and flows. I. The pair correlation function. *J. Phys. A: Math. Gen.*, 19(3):409–422, February 1986. Publisher: IOP Publishing. doi:10.1088/0305-4470/19/3/022.
- 14 Sacha Friedli and Yvan Velenik. Cluster Expansion. In *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*, pages 232–261. Cambridge University Press, Cambridge, 2017. doi:10.1017/9781316882603.006.
- 15 W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57(1):97–109, 1970. Publisher: [Oxford University Press, Biometrika Trust]. doi:10.2307/2334940.
- 16 Tyler Helmuth, Will Perkins, and Guus Regts. Algorithmic Pirogov-Sinai theory. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, pages 1009–1020, New York, NY, USA, June 2019. Association for Computing Machinery. doi:10.1145/3313276.3316305.

- 17 Dmitry Ioffe and Roberto H. Schonmann. Dobrushin-kotecký-Shlosman Theorem up to the Critical Temperature. *Comm Math Phys*, 199(1):117–167, December 1998. doi:10.1007/s002200050497.
- 18 A. Jadbabaie, Jie Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, June 2003. Conference Name: IEEE Transactions on Automatic Control. doi:10.1109/TAC.2003.812781.
- 19 Matthew Jenssen and Will Perkins. Independent sets in the hypercube revisited. *Journal of the London Mathematical Society*, 102(2):645–669, 2020. doi:10.1112/jlms.12331.
- 20 R. Kotecký and D. Preiss. Cluster expansion for abstract polymer models. *Comm. Math. Phys.*, 103(3):491–498, 1986. Publisher: Springer-Verlag. URL: <http://projecteuclid.org/euclid.cmp/1104114796>.
- 21 Shengkai Li, Bahnisikha Dutta, Sarah Cannon, Joshua J. Daymude, Ram Avinery, Enes Aydin, Andréa W. Richa, Daniel I. Goldman, and Dana Randall. Programming active cohesive granular matter with mechanically induced phase changes. *Science Advances*, 7(17):eabe8494, April 2021. Publisher: American Association for the Advancement of Science Section: Research Article. doi:10.1126/sciadv.abe8494.
- 22 Nancy A. Lynch. *Distributed Algorithms*. Elsevier, April 1996.
- 23 Joseph E. Mayer. The Statistical Mechanics of Condensing Systems. I. *J. Chem. Phys.*, 5(1):67–73, January 1937. Publisher: American Institute of Physics. doi:10.1063/1.1749933.
- 24 Joseph E. Mayer. The Theory of Ionic Solutions. *J. Chem. Phys.*, 18(11):1426–1436, November 1950. Publisher: American Institute of Physics. doi:10.1063/1.1747506.
- 25 Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.*, 21(6):1087–1092, June 1953. Publisher: American Institute of Physics. doi:10.1063/1.1699114.
- 26 R. A. Minlos and Ja G. Sinaĭ. The phenomenon of phase separation at low temperatures in some lattice models of a gas. I. *Math. USSR Sb.*, 2(3):335, April 1967. Publisher: IOP Publishing. doi:10.1070/SM1967v002n03ABEH002345.
- 27 Sarah Miracle, Dana Randall, and Amanda Pascoe Streib. Clustering in Interfering Binary Mixtures. In Leslie Ann Goldberg, Klaus Jansen, R. Ravi, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, Lecture Notes in Computer Science, pages 652–663, Berlin, Heidelberg, 2011. Springer. doi:10.1007/978-3-642-22935-0_55.
- 28 R. Olfati-Saber. Flocking for multi-agent dynamic systems: algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, March 2006. Conference Name: IEEE Transactions on Automatic Control. doi:10.1109/TAC.2005.864190.
- 29 G. Ortiz, E. Cobanera, and Z. Nussinov. Dualities and the phase diagram of the p-clock model. *Nuclear Physics B*, 854(3):780–814, January 2012. doi:10.1016/j.nuclphysb.2011.09.012.
- 30 C.-E. Pfister. Interface free energy or surface tension: definition and basic properties. *arXiv*, 2009. arXiv:0911.5232.
- 31 S. A. Pirogov and Ya. G. Sinai. Phase diagrams of classical lattice systems. *Theor Math Phys*, 25(3):1185–1192, December 1975. doi:10.1007/BF01040127.
- 32 S. A. Pirogov and Ya. G. Sinai. Phase diagrams of classical lattice systems continuation. *Theor Math Phys*, 26(1):39–49, January 1976. doi:10.1007/BF01038255.
- 33 Wenquan Qin, Shucong Lin, Xuan Chen, Jian Chen, Lei Wang, Hongpeng Xiong, Qinxi Xie, Zhaohui Sun, Xiujun Wen, and Cai Wang. Food Transport of Red Imported Fire Ants (Hymenoptera: Formicidae) on Vertical Surfaces. *Scientific Reports*, 9(1):3283, March 2019. Number: 1 Publisher: Nature Publishing Group. doi:10.1038/s41598-019-39756-4.
- 34 Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '87, pages 25–34, New York, NY, USA, August 1987. Association for Computing Machinery. doi:10.1145/37401.37406.

- 35 Herbert Robbins. A Remark on Stirling’s Formula. *The American Mathematical Monthly*, 62(1):26–29, 1955. Publisher: Mathematical Association of America. doi:10.2307/2308012.
- 36 Herbert G. Tanner, Ali Jadbabaie, and George J. Pappas. Flocking in Fixed and Switching Networks. *IEEE Transactions on Automatic Control*, 52(5):863–868, May 2007. Conference Name: IEEE Transactions on Automatic Control. doi:10.1109/TAC.2007.895948.
- 37 Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel Type of Phase Transition in a System of Self-Driven Particles. *Phys. Rev. Lett.*, 75(6):1226–1229, August 1995. Publisher: American Physical Society. doi:10.1103/PhysRevLett.75.1226.
- 38 Tamás Vicsek and Anna Zafeiris. Collective motion. *Physics Reports*, 517(3):71–140, August 2012. doi:10.1016/j.physrep.2012.03.004.
- 39 José D. Villa. Swarming Behavior of Honey Bees (Hymenoptera: Apidae) in Southeastern Louisiana. *Ann Entomol Soc Am*, 97(1):111–116, January 2004. Publisher: Oxford Academic. doi:10.1603/0013-8746(2004)097[0111:SBOHBH]2.0.CO;2.
- 40 F. Y. Wu. The Potts model. *Rev. Mod. Phys.*, 54(1):235–268, January 1982. Publisher: American Physical Society. doi:10.1103/RevModPhys.54.235.
- 41 Hai-Tao Zhang, Chao Zhai, and Zhiyong Chen. A General Alignment Repulsion Algorithm for Flocking of Multi-Agent Systems. *IEEE Transactions on Automatic Control*, 56(2):430–435, February 2011. Conference Name: IEEE Transactions on Automatic Control. doi:10.1109/TAC.2010.2089652.

A Details for Aggregation and Alignment in General SOPS

In the general SOPS setting, we can treat the problem as a $q + 1$ -state Potts model on G_Δ with $q + 1$ orientations $\{-1, 0, 1, \dots, q - 1\}$ in which the number of sites assigned -1 is fixed to be exactly $(1 - \rho)N$, where $N = |V(G_\Delta)|$. In other words, sites of the lattice are no longer filled or unfilled, but are instead assigned one of $q + 1$ orientations with the special spin -1 assigned to unoccupied lattice sites. We refer to any edge between particles of differing orientations as “heterogeneous edges,” including those assigned the special orientation -1 .

We again use a Peierls argument to show that for sufficiently large λ , the configuration will compress and one of the q orientations will dominate, with high probability. This proof is an adaptation of the bridging argument used for separation in [5, 6] and thus follows their arguments very closely. The following sections build up to a proof of Theorem 21.

We observe that the result of Theorem 21 will imply both alignment and aggregation (for some values of α and δ) as given in Definitions 10 and 20. The key component of our proof is the construction of a δ -bridge system ($\delta \in (0, 1)$ is a positive constant) for each configuration in $\tilde{\Omega}^{\rho N}$. Recall that a bridge system is a connected network of the long contours of a configuration σ , that is used to “remove” long contours in the Peierls argument to show that they are unlikely. It will also be used to define the region \mathcal{R} required for Theorem 21.

Let E_{wrap} be the set of edges on G_\circlearrowleft corresponding to the edges on G_Δ that wrap around the torus. Thus $|E_{\text{wrap}}| = 2\sqrt{N} - 1$. In a setting with more than three possible orientations, regions of differing orientations are divided up by networks of contours rather than closed walks separating two different orientations. We call these contour networks *complex contours*. Formally, a complex contour refers to a connected subgraph of G_\circlearrowleft of minimum degree at least 2. For a given configuration $\sigma \in \tilde{\Omega}^{\rho N}$, the set of edges \mathcal{C} on G_\circlearrowleft corresponding to its heterogeneous edges will be a union of complex contours. The complex contours of σ thus refers to the edge sets of connected components of the subgraph induced by \mathcal{C} in G_\circlearrowleft .

We now define a bridge system (B, I, Θ) where the set I represents the complex contours in the bridge system, B represents the bridges used to connect these complex contours, and Θ is a mapping that assigns an orientation to each of the components formed after removing the edges of G_Δ corresponding to the edges in I .

► **Definition 22** (Bridge Systems). Fix $\delta > 0$. Consider a tuple (B, I, Θ) , where B and I are subsets of $E(G_\square)$ and $\Theta : V(G_\Delta) \rightarrow \{-1, 0, 1, \dots, q-1\}$ is a function assigning each vertex an orientation or the value -1 (which we will use to represent vacant sites). We say (B, I, Θ) is a δ -bridge system if:

1. The subgraph induced in G_\square by I has no vertex of degree less than 2. Practically, I represents a union of complex contours that subdivides G_Δ into regions.
2. The subgraph induced in G_\square by $B \cup I \cup E_{\text{wrap}}$ is connected and has no vertex of degree less than 2.
3. $B \cap I = \emptyset$ and $|B| \leq \frac{1-\delta}{2\delta}|I|$
4. For any two neighboring sites $u, v \in G_\Delta$, $\Theta(u) = \Theta(v)$ if and only if the dual edge corresponding to $\{u, v\}$ is not in I .

Consider a set of edges I , that is a union of the edge sets of complex contours. Let σ be a configuration in $\tilde{\Omega}^{\rho N}$. We say a complex contour C of σ is *bridged* (by I) if $C \subseteq I$. We say a site v is bridged (by I) if there is a path over G_Δ using only sites of the same orientation (including -1) in σ as v to a site incident to an edge in I . Consider a region $R \subseteq V(G_\Delta)$ that is connected as an induced subgraph of G_Δ . We call R a *bridged region* if $bd(R) \subseteq I$ and a *minimal bridged region* if there is no bridged region R' where $R' \subsetneq R$. Notably, the edge set I partitions $V(G_\Delta)$ into minimal bridged regions.

► **Definition 23** (Bridge System for a Configuration). Fix $\delta > 0$ and a configuration $\sigma \in \tilde{\Omega}^{\rho N}$. We say a tuple (B, I, Θ) is a δ -bridge system for a configuration σ if

1. Each minimal bridged region R by (B, I, Θ) contains at most $\delta|R|$ unbridged particles.
2. No complex contour C of σ meets any edge in $B \cup I \cup E_{\text{wrap}}$. Formally, the edge-induced subgraphs $G_\Delta[C]$ and $G_\Delta[B \cup I \cup E_{\text{wrap}}]$ do not share any vertices.
3. For each minimal bridged region R , $\Theta(v)$ must have the same value for every site $v \in R$ and this value $\Theta(v)$ must correspond to the orientation in σ of some bridged particle in R .

► **Definition 24** (Orientation of a Minimal Bridged Region). Given a δ -bridge system (B, I, Θ) for a configuration $\sigma \in \tilde{\Omega}^{\rho N}$. We can associate with each minimal bridged region R of I an orientation $y_R \in \{-1, 0, 1, \dots, q-1\}$.

To determine y_R , we denote by R^* the set of sites $v \in R$ with a path over G_Δ using only sites of the same orientation in σ as v to a site incident to an edge in $bd(R)$. We note that $bd(R) \subseteq I$ and the edges $B \cup I \cup E_{\text{wrap}}$ connect the components of $bd(R)$ in G_\square . This implies that every vertex in R^* must have the same orientation in σ , as any contour C between regions of differing orientations in R^* must intersect $B \cup I \cup E_{\text{wrap}}$, implying that C also must be included in the set I , allowing us to subdivide R , contradicting its minimality. The orientation y_R of R is thus defined to be the common orientation of the sites of R^* .

Thus, for each minimal bridged region R with orientation y_R , we must have $\Theta(v) = y_R$ for all $v \in R$. The proofs of the Lemmas will be given in the long version of the paper.

Our next step is to associate with each $\sigma \in \tilde{\Omega}^{\rho N}$ a δ -bridge system.

► **Lemma 25.** For each $\sigma \in \tilde{\Omega}^{\rho N}$ and $\delta \in (0, 1)$, we can construct a δ -bridge system $\mathcal{B}_\delta(\sigma) = (B_\delta(\sigma), I_\delta(\sigma), \Theta_\delta(\sigma))$.

Without reference to any specific configuration in $\tilde{\Omega}^{\rho N}$, we use the connectedness requirement of bridge systems to compute an upper bound on the number of bridge systems that is exponential on $|I|$. This is important as the Peierls argument “removes” the heterogeneous edges in I , which gives an improvement in weight of a similar order of growth.

► **Lemma 26.** *The number of δ -bridge systems (B, I, Θ) where $|I| = \ell$ is at most $7 \cdot 6^{2\sqrt{N}-1} \cdot (3(q+1))^{\frac{1+\delta}{2\delta}\ell}$.*

Assuming $\delta \in (0, \rho)$, we define $\tilde{\Omega}_\ell^{\rho N} := \{\sigma \in \tilde{\Omega}^{\rho N} : |I_\delta(\sigma)| = \ell\}$, where $I_\delta(\sigma)$ is comes from the δ -bridge system constructed for σ . Also, let $\tilde{\Omega}^{\leq \delta N}$ be the the set of configurations over G_Δ where at least $(1-\delta)N$ sites have orientation -1 (this corresponds to empty sites in our model). Note that $\tilde{\Omega}^{\leq \delta N} \not\subseteq \tilde{\Omega}^{\rho N}$. For the Peierls argument, we define two functions, $f_\ell^1 : \tilde{\Omega}_\ell^{\rho N} \rightarrow \tilde{\Omega}^{\leq \delta N}$ and $f^2 : \tilde{\Omega}^{\leq \delta N} \rightarrow \tilde{\Omega}^{\rho N}$. The function f_ℓ^1 is used to erase the heterogeneous edges in I , creating a configuration of significantly higher weight, though not one with ρN particles. To fix this, a second function, f^2 is used to restore the number of particles back to ρN . This way, $f^2 \circ f_\ell^1$ maps each σ in $\tilde{\Omega}_\ell^{\rho N}$ to a valid configuration with exactly ρN filled sites. The definitions of f_ℓ^1 and f^2 are given in the full version of the paper.

As the bridge system with just a polynomial amount of additional information can be used to reconstruct σ from $f^2 \circ f_\ell^1$, our upper bound on the number of bridge systems can be used to upper bound $|(f^2 \circ f_\ell^1)^{-1}(\tau)|$ for any τ in the image of $f^2 \circ f_\ell^1$. This allows us to prove the following Lemma:

► **Lemma 27.** *Fix $\rho < \frac{1}{3}$, any $\alpha > 1$, $\delta \in (0, \min\{\rho, 1 - \frac{1}{\alpha^2}\})$ and $\lambda > \lambda_0(q, \rho, \alpha, \delta)$ sufficiently large, where:*

$$\lambda_0(q, \rho, \alpha, \delta) := \left((3(q+1))^{\alpha \frac{1+\delta}{2\delta}} 36^{\frac{1}{4\sqrt{3\rho}}} \right)^{\frac{1}{\alpha - \frac{1}{\sqrt{1-\delta}}}}.$$

Denote by $\tilde{\Omega}_{\geq \alpha \cdot bd_{\min}(\rho N)}^{\rho N}$ the set of configurations σ where $|I_\delta(\sigma)| \geq \alpha \cdot bd_{\min}(\rho N)$, where $bd_{\min}(k)$ is the minimum possible boundary length of a region of $k \in \mathbb{N}$ particles. Then there exists a constant $\zeta = \zeta(q, \rho, \alpha, \delta, \lambda) < 1$ such that $\tilde{\pi}_{\text{Potts}}(\tilde{\Omega}_{\geq \alpha \cdot bd_{\min}(\rho N)}^{\rho N}) < \zeta^{\sqrt{N}}$ for all sufficiently large values of N .

As the bridge system with just a polynomial amount of additional information can be used to reconstruct σ from $f^2 \circ f_\ell^1$, our upper bound on the number of bridge systems can be used to upper bound $|(f^2 \circ f_\ell^1)^{-1}(\tau)|$ for any τ in the image of $f^2 \circ f_\ell^1$. This allows us to prove the following Lemma:

The use of Lemma 27 along with some results on the minimum possible boundary lengths of regions of k particles allows us to show that there will exist a low perimeter region dominated by a single color, allowing us to prove Theorem 21.