

A Markov Chain Algorithm for Compression in Self-Organizing Particle Systems *

Sarah Cannon[†]
Georgia Institute of Technology
Atlanta, GA, USA
sarah.cannon@gatech.edu

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

Joshua J. Daymude[‡]
Arizona State University
Tempe, AZ, USA
jdaymude@asu.edu

Andréa W. Richa[‡]
Arizona State University
Tempe, AZ, USA
aricha@asu.edu

ABSTRACT

We consider programmable matter as a collection of simple computational elements (or *particles*) with limited (constant-size) memory that self-organize to solve system-wide problems of movement, configuration, and coordination. Here, we focus on the *compression* problem, in which the particle system gathers as tightly together as possible, as in a sphere or its equivalent in the presence of some underlying geometry. More specifically, we seek fully distributed, local, and asynchronous algorithms that lead the system to converge to a configuration with small perimeter. We present a Markov chain based algorithm that solves the compression problem under the *geometric amoebot model*, for particle systems that begin in a connected configuration with no holes. The algorithm takes as input a bias parameter λ , where $\lambda > 1$ corresponds to particles favoring inducing more lattice triangles within the particle system. We show that for all $\lambda > 5$, there is a constant $\alpha > 1$ such that at stationarity with all but exponentially small probability the particles are α -compressed, meaning the perimeter of the system configuration is at most $\alpha \cdot p_{min}$, where p_{min} is the minimum possible perimeter of the particle system. We additionally prove that the same algorithm can be used for *expansion* for small values of λ ; in particular, for all $0 < \lambda < \sqrt{2}$, there is a constant $\beta < 1$ such that at stationarity, with all but an ex-

ponentially small probability, the perimeter will be at least $\beta \cdot p_{max}$, where p_{max} is the maximum possible perimeter.

Keywords

Self-organizing Particles; Compression; Markov Chains

1. INTRODUCTION

Many programmable matter systems have recently been proposed and realized—modular and swarm robotics, synthetic biology, DNA tiling, and smart materials form an incomplete list—and each is often tailored toward a specific task or physical setting. In our work on self-organizing particle systems, we abstract away from specific settings and instead describe programmable matter as a collection of simple computational elements (to be referred to as *particles*) with limited computational power that each perform fully distributed, local, asynchronous algorithms to solve system-wide problems such as movement, configuration, and coordination. Here we present an algorithm for *compression*, in which the particle system gathers as tightly together as possible, as in a sphere or its equivalent in the presence of some underlying geometry. This phenomenon is often found in natural systems: fire ants form floating rafts by gathering in such a manner, and honey bees communicate foraging patterns within their hives. While each individual ant or bee cannot view the group as a whole when soliciting information, it can take cues from its immediate neighbors to achieve cooperation. It is with this motivation that we present a distributed algorithm for compression in the amoebot model derived from a Markov chain process.

In the (geometric) amoebot model, more formally defined in Section 2.1, particles with limited computational power move among the vertices of the triangular lattice Γ (Figure 1(a)) by traveling along the edges of Γ . The *compression problem* seeks to reorganize the configuration of a particle system (via movements of particles) such that the system converges to a configuration with small perimeter, where we measure the perimeter of a configuration by the length of the walk along its boundary. We say a particle system is α -compressed, for $\alpha > 1$, if the perimeter of the particle configuration is at most α times the minimum possible perimeter for those particles.

* A full version of this paper, including omitted proofs, is available at www.arxiv.org/abs/1603.07991.

[†]Supported in part by NSF DGE-1148903 and a grant from the Simons Foundation (#361047 to Sarah Cannon)

[‡]Supported in part by NSF awards CCF-1353089, CCF-1422603, and REU-026935.

[§]Supported in part by NSF grant CCF-1526900.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'16, July 25–28, 2016, Chicago, IL, USA

© 2016 ACM. ISBN 978-1-4503-3964-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2933057.2933107>

1.1 Our results and techniques

We present a *Markov chain* $\mathcal{M}$ for particle compression under the geometric amoebot model that can be directly translated into a *fully distributed, local, asynchronous compression algorithm* $\mathcal{A}$. Both $\mathcal{A}$ and $\mathcal{M}$ take as input a *bias parameter* λ (where $\lambda > 1$ makes induced lattice triangles more favorable) and start from an arbitrary initial configuration for the particles that is connected and has no holes.

Markov chain $\mathcal{M}$ is carefully designed according to the distributed and local nature of the system, so that the particles always stay connected and no holes form. Furthermore, we prove $\mathcal{M}$ is *reversible and ergodic*, meaning many of the standard tools of Markov chain analysis can be applied. While most of these proofs rely only on first principles, we emphasize they are far from trivial; working in a distributed setting necessitated carefully defined protocols for local moves that made proofs challenging.¹

When the particles execute the local moves of $\mathcal{M}$ (by running $\mathcal{A}$) for long enough, the configuration of the particles converges to the stationary distribution of $\mathcal{M}$. We prove for all large enough λ there is a constant $\alpha = \alpha(\lambda) > 1$ such that at stationarity, with all but exponentially small probability, the particles are α -compressed, meaning the perimeter of the particle configuration is at most α times the minimum perimeter (which is $\Theta(\sqrt{n})$ for systems of n particles). We additionally show the counterintuitive result that $\lambda > 1$ is not enough to guarantee compression. In fact, for all $0 < \lambda < \sqrt{2}$, there is a constant $\beta < 1$ such that at stationarity with all but exponentially small probability the perimeter is at least a β fraction of the maximum perimeter, which is $\Theta(n)$ for systems of n particles. We call such a configuration β -expanded. This implies that for any $0 < \lambda < \sqrt{2}$, the probability that the particles are α -compressed is exponentially small for any constant α .

The motivation underlying the design of this Markov chain is from statistical physics, where ensembles of particles reminiscent of our amoebot model are used to study physical systems. Like a spring relaxing, systems tend to favor configurations that minimize energy. The energy function is determined by a *Hamiltonian* $H(\sigma)$; each configuration σ has weight $w(\sigma) = e^{-B \cdot H(\sigma)} / Z$, where $B = 1/T$ is inverse temperature and $Z = \sum_{\tau} e^{-B \cdot H(\tau)}$ is the normalizing constant known as the *partition function*.

In our amoebot model, we assign each configuration σ a Hamiltonian $H(\sigma) = -t(\sigma)$, where $t(\sigma)$ is the number of triangles in σ , i.e., the number of faces of the triangle lattice Γ with all three vertices occupied by particles. Setting $\lambda = e^B$, we get $w(\sigma) = \lambda^{t(\sigma)}$. As λ gets larger (by increasing B , effectively lowering temperature), we favor configurations with a large number of occupied triangles, causing increasingly compressed configurations. Favoring edges with both endpoints occupied is an alternative metric we could consider, but we prove for connected configurations on Γ without holes the two measures are equivalent. Likewise, favoring shorter perimeter is a third equivalent representation.

The key tool used to establish compression is a careful *Peierls argument*, used in statistical physics to study non-uniqueness of limiting Gibbs measures and determining the presence of phase transitions (see, e.g., [9]), and in computer science to establish slow mixing of Markov chains (see, e.g., [2]). For standard Peierls arguments, configura-

tions typically are not required to be connected and can have holes. Here, we focus on the connected and hole-free model to show compression can be achieved, even in an asynchronous distributed system where particles have constant-size memory and use only local information.

1.2 Related Work

Our algorithm is derived from a carefully designed Markov chain, enabling us to provide provable guarantees of its behavior. Random particle processes on the grid with hard constraints (e.g., *simple exclusion processes* where no two particles can occupy the same location) have been studied in statistical physics, but we want results under the further constraints of distributed computing. Our work exploits the memoryless, stochastic nature of Markov chains to accomplish particle compression in the amoebot model, just one of many distributed physical systems where compression-type problems have been studied.

When considering physical systems and models, one can differentiate between active and passive systems. Particles in passive systems have no explicit control over their movements, and in some cases do not have any computational power. One example is the DNA self-assembly work described in [19], where strands of DNA gather together to form larger structures with certain patterns.

In active systems, particles have control over their behavior and—depending on the model—can achieve some directed locomotion. *Swarm robotics* is one example; different variations of shape formation and collection problems have been studied (e.g. [11, 17]), but always with more computational power or global knowledge of the system. Similarly, pattern formation and creation of convex structures has been studied in the *cellular automata* domain (e.g. [4, 8]), but differs from our model by assuming more powerful computational capabilities. The *nubot* model [20] addresses a framework for biomolecular-inspired models which—although allowing some non-local movements—provides additional ways to create two dimensional shapes in polylogarithmic time.

Nature offers a variety of examples in which gathering and cooperative behavior is apparent. For example, social insects often exhibit compression-like characteristics in their collective behavior: fire ants form floating rafts [15], cockroach larvae perform self-organizing aggregation [13, 16], and honey bees choose hive locations based on a decentralized process of swarming and recruitment [3].

The rendezvous (or gathering) problem seeks to gather mobile agents together on some node of a graph (see e.g. [1] and the references within). In comparison, our particles follow the exclusion principle, and hence would not be able to gather at a single node, and are computationally simpler than the mobile agents considered.

Lastly, in [5, 6], algorithms for hexagon shape formation in the amoebot model were presented. Although a hexagon satisfies a compressed configuration as we define here, the Markov chain-based algorithm we present takes a fully decentralized and local approach. This is naturally self-stabilizing, forgoing the need for a seed particle that may coordinate or initiate some underlying organization of the set of particles, as required in [5] and even more critically in [6].

¹Due to length constraints some proofs were omitted.




Figure 1: (a) A section of the triangular lattice Γ ; (b) expanded particles (each denoted by its two occupied adjacent locations in Γ and a thick line in between) and contracted particles (occupying one location).

2. BACKGROUND AND MODEL

We begin with the geometric amoebot model for programmable matter. We then define some properties of particle systems and discuss what it means for a particle system to be compressed. We conclude with an overview of Markov chains, which form the basis for our algorithm for compression to be presented in Section 3.

2.1 The Amoebot Model

In the *amoebot model* [7], programmable matter consists of particles whose properties we now detail. An infinite undirected graph $G = (V, E)$ represents the set of relative locations that the particles can occupy (given by V) and the set of all possible atomic transitions between locations in V (given by E) [7]. We further assume the *geometric* variant of the amoebot model, which imposes an underlying geometric structure $G = \Gamma$, where Γ is the *triangular lattice* shown in Figure 1(a) (also called the infinite regular triangular grid graph, and denoted by G_{eqt} in earlier work).

Each particle occupies either a single location (i.e., it is *contracted*) or a pair of two adjacent locations on the graph (i.e., it is *expanded*); Figure 1(b) illustrates expanded and contracted particles on Γ . Each location can be occupied by at most one particle. Particles achieve movement via a series of *expansions* and *contractions*: a contracted particle may expand into an adjacent unoccupied location to become expanded, and completes its movement by contracting to once again occupy only one location.

Two particles occupying adjacent nodes are said to be connected by a bond (and hence induce an *edge* in the particle system), and we refer to them as *neighbors*. Particles are *anonymous*, but can uniquely identify each one of their (six) possible bonds and check which bonds lead to nodes occupied by neighboring particles. Additionally, the particles do not share any global orientation or coordinate system.

Every particle has a constant-size, shared, local memory which both it and its neighbors can read and write to for communication. Because of the limitation on memory size, particles cannot know neither the total size of the system nor an estimate of it. Particles execute a sequence of *atomic actions*, in each of which they do some local computation (in our case, those may involve checking which of its adjacent locations are occupied with particles) and an expansion or contraction. We assume a fully asynchronous system, where particles perform atomic actions concurrently and at different, possibly variable speeds; conflicts, which in our context arise when two particles attempt to expand into the same location, are resolved in an arbitrary manner. In order to analyze such systems, we use the standard asynchronous model from distributed computing, allowing us to evaluate the progress of the system through a sequential series of individual *particle activations*, where every time a particle is activated, it performs an atomic action.

2.2 Terminology for Particle Systems

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

We will let capital letters refer to particles and lower case letters refer to locations on the triangular lattice Γ , e.g., “particle P at location ℓ .” For a particle P (resp., location ℓ), we use $N(P)$ (resp., $N(\ell)$) to denote the set of particles adjacent to P (resp., to ℓ), where by *adjacent* we mean connected by a lattice edge. Similarly, for a particle P (resp., location ℓ), we will use $n(P)$ (resp., $n(\ell)$) to denote the six locations in the neighborhood of P (resp., of ℓ). For locations ℓ and ℓ' , by $N(\ell \cup \ell')$ we mean $(N(\ell) \cup N(\ell')) \setminus \{\ell, \ell'\}$; the same holds for $n(\ell \cup \ell')$.

By an *edge* or *triangle* of a configuration σ we mean, respectively, an edge or face of Γ such that all (respectively, two or three) incident vertices are occupied by particles. The number of edges of σ is $e(\sigma)$ and the number of triangles is $t(\sigma)$. Throughout, by a *path* or a *cycle* we mean a path or cycle in the underlying graph Γ where all vertices are occupied by particles, and in the case of a cycle, at least one location inside the cycle is unoccupied. Two particles are *connected* if there exists a path between them, and a configuration is *connected* if all pairs of particles are. A *hole* in a configuration is a maximal finite component of adjacent unoccupied locations. We specifically consider connected configurations with no holes, and our algorithm, if starting at such a configuration, will maintain these properties.

2.3 Compression of Particle Systems

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

For a connected configuration σ of n particles with no holes, the *perimeter* of σ , denoted $p(\sigma)$, is the length of the walk around the (single external) boundary of the particles. In an abuse of notation, we use the term *perimeter* to refer both to the length $p(\sigma)$ of this walk and the walk itself. We assume any walk $\mathcal{W}$ along the perimeter of a configuration is in the clockwise direction. For a walk $\mathcal{W}$ in Γ with some notion of direction, we say location ℓ is *left* of edge $e = \overrightarrow{st} \in \mathcal{W}$ traversed from s to t if ℓ , s and t are the vertices of a triangular face of Γ with ℓ the next vertex counterclockwise around this face from t . With this terminology, every edge traversed in a clockwise walk $\mathcal{W}$ along the perimeter of σ has an unoc-

cupied location to its left. Specifically, for any consecutive particles A, B, C in $\mathcal{W}$, the locations in the clockwise span of $n(B)$ from A to C are always unoccupied and there is at least one of them. Note an edge may appear twice in a perimeter walk $\mathcal{W}$; in this case, its length is counted twice in $p(\sigma)$. For a connected configuration of n particles without holes, the perimeter ranges from a maximum value of $2n - 2$ when the particles are in their least compressed state (a tree with no induced triangles) to some minimum value $p_{\min}(n) = \Theta(\sqrt{n})$ when the particles are in their most compressed state. It is easy to see $p_{\min}(n) \leq 4\sqrt{n}$, and we now prove any configuration σ of n particles has $p(\sigma) \geq \sqrt{n}$; this bound is not tight but suffices for our proofs.

LEMMA 2.1. *A connected configuration with $n \geq 2$ particles and no holes has perimeter at least $\sqrt{n}$.*

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

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

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

If Q has two neighbors, removing Q from σ produces two connected particle configurations σ_1 and σ_2 , where σ_1 has n_1 particles, σ_2 has n_2 particles, and $n_1 + n_2 = n - 1$. Thus

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

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

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

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

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

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

When n is clear from context we omit it and refer to $p_{\min} = p_{\min}(n)$ and $p_{\max} = p_{\max}(n)$. We now define what it means for a particle system to be compressed.

DEFINITION 2.2. *For any $\alpha > 1$, a connected configuration σ with no holes is α -compressed if $p(\sigma) \leq \alpha \cdot p_{\min}$.*

We prove in Section 4 that our algorithm, when executed for a sufficiently long time, achieves α -compression with all but exponentially small probability for any constant $\alpha > 1$, provided n is sufficiently large. We note α -compression

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

In order to minimize perimeter using only simple local moves, we exploit the following relationship.

LEMMA 2.3. *For a connected particle configuration σ with no holes, $t(\sigma) = 2n - p(\sigma) - 2$.*

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

Minimizing perimeter is also equivalent to maximizing edges, as we now show.

LEMMA 2.4. *For a connected particle configuration σ with no holes, $e(\sigma) = 3n - p(\sigma) - 3$.*

PROOF. The proof is nearly identical to that of Lemma 2.3, counting particle-edge incidences instead. $\square$

COROLLARY 2.5. *For a connected particle configuration σ with no holes, $t(\sigma) = e(\sigma) - (n - 1)$.*

COROLLARY 2.6. *A connected particle configuration σ with no holes and minimum perimeter is also a configuration with the maximum number of triangles (and edges).*

2.4 Markov chains

The distributed protocol for particle compression we present is based on a Markov chain, i.e., a memoryless stochastic process defined on a finite set of states Ω . The transition matrix P on $\Omega \times \Omega \rightarrow [0, 1]$ is defined so that $P(x, y)$ is the probability of moving from state x to state y in one step, for any pair $x, y \in \Omega$. The t -step transition probability $P^t(x, y)$ is the probability of moving from x to y in exactly t steps.

A Markov chain is *ergodic* if it is *irreducible*, i.e., for all $x, y \in \Omega$, there is a t such that $P^t(x, y) > 0$, and *aperiodic*, i.e., for all $x, y \in \Omega$, $\text{g.c.d.}\{t : P^t(x, y) > 0\} = 1$. Any finite, ergodic Markov chain converges to a unique distribution π , i.e., for all $x, y \in \Omega$, $\lim_{t \rightarrow \infty} P^t(x, y) = \pi(y)$. In fact, for any distribution π' such that $\pi'(x)P(x, y) = \pi'(y)P(y, x)$ (the *detailed balance condition*) and $\sum_{x \in \Omega} \pi'(x) = 1$, π' is the unique stationary distribution of $\mathcal{M}$ (see, e.g., [10]).

Given a desired stationary distribution π on Ω , the celebrated Metropolis-Hastings algorithm [12] defines appropriate transition probabilities. Starting at state x , pick a neighbor y in Ω uniformly with probability $1/(2\Delta)$, where Δ is the maximum number of neighbors of any state, and move to y with probability $\min(1, \pi(y)/\pi(x))$; with the remaining probability stay at x and repeat. Using detailed balance, one can verify if the state space is connected then π must be the stationary distribution. While calculating $\pi(x)/\pi(y)$ seems to require global knowledge, this ratio can often be calculated using only local information when many terms cancel

out. In our case, the Metropolis probabilities are simply $\min(1, \lambda^{|t(y)-t(x)|})$; this difference is just the change in the number of triangles incident to the moving particle, which can be calculated with only local information.

3. ALGORITHMS FOR COMPRESSION

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

Our algorithm carefully maintains several critical properties throughout its execution. First, the particle system stays connected and no holes form, even while particles decide where to move based only on local information. Additionally, any moves made are reversible: if a particle moves to a new location, there is a nonzero probability that during the next step, it moves back to its previous location. Finally, the moves allowed by the algorithm (and respective Markov chain) suffice to transform any configuration of particles into any other configuration. These conditions are essential so certain tools from Markov chain analysis can be applied.

In addition to the precise conditions needed to ensure connectivity and reversibility, our algorithm achieves compression by making particles more likely to move into a position where they form more triangles with their neighbors. Specifically, a bias parameter λ controls how strongly the particles favor being incident to triangles; $\lambda > 1$ corresponds to favoring triangles, while $\lambda < 1$ corresponds to disfavoring triangles. As Lemma 2.3 shows, locally favoring more triangles is equivalent to globally favoring a shorter perimeter; this is the relationship we exploit to obtain particle compression.

3.1 The Local Algorithm $\mathcal{A}$

We start with two key properties that enable a particle to move from location ℓ to location ℓ' . If ℓ and ℓ' are neighboring locations on Γ , let $\mathbb{S} = N(\ell) \cap N(\ell')$ be the set of particles adjacent to both ℓ and ℓ' (i.e., $|\mathbb{S}| = 0, 1$, or 2).

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

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

These properties capture precisely the structure required to maintain particle connectivity and prevent holes from forming. Additionally, both are symmetric for ℓ and ℓ' , necessary for reversibility. However, they are not so restrictive as to limit the movement of particles and prevent compression from occurring; we will see moves satisfying these properties suffice to transform any configuration into any other.

We now present the local asynchronous algorithm that each particle runs. Later we show how to view this algorithm through the lens of our Markov chain $\mathcal{M}$. Both $\mathcal{A}$ and $\mathcal{M}$ take as an input a bias parameter $\lambda > 1$, and begin in an

arbitrary starting configuration σ_0 of contracted particles that is connected and has no holes.

Algorithm $\mathcal{A}$: Algorithm for compression run by particle P

If P is contracted:

- 1: Let ℓ denote P 's current location.
- 2: Particle P chooses neighboring location ℓ' uniformly at random from the six possible choices, and generates a random number $q \in (0, 1)$.
- 3: **if** ℓ' is unoccupied and P has no expanded neighboring particle at ℓ **then**
- 4: P expands to simultaneously occupy ℓ and ℓ' .
- 5: **if** there are no expanded particles adj. to ℓ or ℓ' **then**
- 6: P sets $flag = TRUE$ in its local memory
- 7: **else** P sets $flag = FALSE$.

If P is expanded:

- 8: Let t be the number of triangles formed by P when it was contracted at its original location ℓ , and let t' be the number of triangles formed by P if it contracts at the location it expanded into, ℓ' .
 - 9: **if** (1) location ℓ does not have five neighboring particles, (2) locations ℓ and ℓ' satisfy Property 1 or Property 2, (3) $q < \lambda^{t'-t}$, and (4) $flag = TRUE$ **then**
 - 10: P contracts to ℓ' .
 - 11: **else** P contracts back to ℓ .
-

Each particle P continuously runs Algorithm $\mathcal{A}$, executing Steps 1–7 if P is contracted, and Steps 8–11 if P is expanded. Note a constant number of random bits suffice to generate q , as only a constant precision is required (given that $t' - t$ is an integer in $[-6, 6]$ and λ is a constant). In Step 9, Condition (1) ensures no holes form; Condition (2) ensures the particles stay connected and the corresponding Markov chain $\mathcal{M}$ is reversible; Condition (3) ensures the moves happen with probabilities such that $\mathcal{M}$ converges to the desired distribution; and Condition (4) ensures P is the only particle in its neighborhood potentially moving to a new position, ℓ' , since P last expanded. The claims regarding Conditions (1)–(3) pertain to the underlying Markov process and will be formalized in the next subsections. To understand why Condition (4) suffices to ensure P is the only particle (potentially) moving to a new location in its immediate neighborhood, note that once P expands, any other particle P' expanding into a position adjacent to P will contract back to its original position, since P' will set its own flag to FALSE in Step 7. We assume any conflicts arising from two particles attempting to concurrently move into the same location are resolved arbitrarily. Hence, any concurrent movements will cover pairwise disjoint neighborhoods and the respective actions will be mutually independent.

Following the classical asynchronous model [14], for any starting configuration σ_0 , we have that for any concurrent execution of $\mathcal{A}$ that reaches a configuration σ , there is a sequential execution of the same atomic actions that also reaches σ , for all σ .

3.2 The Markov Chain $\mathcal{M}$

We now make explicit the Markov chain $\mathcal{M}$ corresponding to Algorithm $\mathcal{A}$. The state space of $\mathcal{M}$ is the set of all connected configurations of n contracted particles and no holes.

Transitions between states happen subject to the rules and probabilities we now present.

Beginning at any connected, hole-free configuration σ_0 of contracted particles, repeat:

1. Select particle P uniformly at random from among all particles; let ℓ be its location. Choose, uniformly at random, $\ell' \in n(\ell)$ and $q \in (0, 1)$.
2. **If** ℓ' is unoccupied, **then** P expands to simultaneously occupy ℓ and ℓ' ; **Else** return to Step 1.
3. Let t be the number of triangles formed by P in position ℓ , and let t' be the number of triangles formed by P in position ℓ' .
4. **If** (1) location ℓ does not have five neighboring particles, (2) locations ℓ and ℓ' satisfy Property 1 or Property 2, and (3) $q < \lambda^{t'-t}$ **then** contract P to ℓ' .
5. **Else** contract P back to ℓ .

In $\mathcal{M}$, we have paired up the consecutive actions that a particle will take according to Algorithm $\mathcal{A}$ when contracted and then expanded into one state transition, so that a transition corresponds to the complete movement of exactly one (contracted) particle into a new (adjacent) location. Let P be a particle that eventually moves from location ℓ to ℓ' at some time t' according to an execution of $\mathcal{A}$. Condition (4) in Step 9 and the condition in Step 3 of $\mathcal{A}$ ensure that if P was activated and expanded at time t before being activated once again and contracting to occupy ℓ' at time $t' > t$, even if other particles are activated between t and t' no movement to a new location will occur in the neighborhood of P in that time interval. Hence all activations of other particles in the neighborhood of P in the interval (t, t') can be ignored, justifying the pairing of actions in the Markov chain $\mathcal{M}$.

For every sequential execution of atomic actions that leads to configuration σ' in $\mathcal{A}$, there exists a sequence of transitions in $\mathcal{M}$ that reaches a configuration σ such that σ can be obtained from σ' by preserving the locations of all contracted particles in σ' and by letting every expanded particle in σ' contract according to the rules of $\mathcal{A}$, for all corresponding pairs (σ', σ) . Conversely, every sequence of transitions in $\mathcal{M}$ that reaches a configuration σ directly corresponds to a sequence of atomic actions in $\mathcal{A}$ also leading to $\sigma' = \sigma$. The perimeter of the respective σ' and σ differs by at most a constant factor, and hence proving α -compression for σ also implies α' -compression for σ' , and vice-versa, for constants α and α' . Hence, we can use $\mathcal{M}$, and respective Markov chain tools and techniques, in order to analyze the correctness of algorithm $\mathcal{A}$.

We note that, under the assumptions of the asynchronous model of distributed computing, one cannot typically assume the next particle to be activated is equally likely to be any particle, as we assume in Step 1 of the description of $\mathcal{M}$. We make this assumption in order to be able to explicitly calculate the stationary distribution of $\mathcal{M}$ so that we can provide rigorous guarantees about its structure, but do not expect the behavior of the system would be substantially different if this requirement was relaxed.

To justify this random activation assumption, we note that random sequences of particle activations can be ap-

proximated using Poisson clocks with mean 1.² That is, each particle activates, and executes Algorithm $\mathcal{A}$, at a random real time drawn from the exponential distribution e^{-t} . After each action, the particle then computes another random time drawn from the same distribution e^{-t} and executes again after that amount of time has elapsed. The exponential distribution is unique in that, if particle P has just activated, it is equally likely that any particle will be the next particle to activate, including particle P (see, e.g., [10]). Moreover, the particles update without requiring knowledge of any of the other particles' clocks. Similar Poisson clocks are commonly used to describe physical systems that perform updates in parallel in continuous time.

3.3 Invariants for Markov chain $\mathcal{M}$

We begin by showing $\mathcal{M}$ maintains certain invariants. We prove Conditions (1) and (2) in Step 4 of Markov chain $\mathcal{M}$ ensure the particles remain in a connected configuration with no holes, provided they start in such a configuration.

LEMMA 3.1. *If the particles are initially connected, during the execution of Markov chain $\mathcal{M}$ they remain connected.*

PROOF. Consider one iteration of $\mathcal{M}$ where a particle P moves from location ℓ to location ℓ' . Let σ be the configuration before this move, and σ' the configuration after. We show if σ is connected, then so is σ' .

A move of particle P from ℓ to ℓ' occurs only if ℓ and ℓ' are adjacent and satisfy Property 1 or Property 2. First, suppose they satisfy Property 1. Let $P_1, P_2 \neq P$ be particles, and let Q be a path connecting them in σ . If $P \notin Q$, then P_1 and P_2 remain connected by Q in σ' . If $P \in Q$, let N_1 and N_2 be the two particles on path Q before and after P , respectively. By Property 1, there exist paths in $N(\ell \cup \ell')$ from N_1 to a particle $S_1 \in \mathbb{S}$ and from N_2 to a particle $S_2 \in \mathbb{S}$, possibly with $S_1 = S_2$. A (not necessarily simple) path from P_1 to P_2 in σ' is the union of Q from P_1 to N_1 ; the path from N_1 to S_1 in $N(\ell \cup \ell')$; if $S_1 \neq S_2$, the path from S_1 to P at location ℓ' to S_2 ; the path from S_2 to N_2 in $N(\ell \cup \ell')$; and Q from N_2 to P_2 . As P is connected to $S_1 \in \mathbb{S}$ and by the above argument S_1 is connected to all other particles, we conclude P is connected to every other particle and thus σ' is connected.

Next, assume locations ℓ and ℓ' satisfy Property 2. Let $P_1, P_2 \neq P$ be particles; we show they are connected by a path not containing P . Path Q connecting P_1 and P_2 in σ exists, and suppose it contains P . Let N_1 and N_2 be the vertices on Q before and after P , respectively. Both N_1 and N_2 are neighbors of ℓ , and by Property 2 all neighbors of ℓ are connected by a path in $N(\ell)$. Thus Q can be augmented to form a (not necessarily simple) walk Q' by replacing P with a path from N_1 to N_2 in $N(\ell)$. As $P \notin Q'$, this walk also connects P_1 and P_2 in σ' . Additionally, because ℓ' has at least one neighbor by Property 2, P remains connected to all other particles in σ' and thus σ' is connected. $\square$

LEMMA 3.2. *If the particles begin in a connected configuration with no holes, during the execution of Markov chain $\mathcal{M}$ they will never form a configuration with a hole.*

PROOF. Recall we assume a cycle in σ encircles at least one unoccupied location; note a configuration has a hole if

²The analysis can be modified to accommodate each clock having its own constant mean; however, for the sake of ease of presentation, we assume here that they are all i.i.d.

and only if it has a cycle encircling that hole. Let σ be a particle configuration with particle P at location ℓ , and σ' the same configuration with P at neighboring location ℓ' . We assume σ has no cycles (i.e., no holes) and prove σ' has no cycles (i.e., no holes).

We first show if a cycle is introduced in σ' , then P must be on that cycle. Suppose this is not the case and σ' has a cycle $\mathcal{C}$ with $P \notin \mathcal{C}$. If P is removed from location ℓ' , $\mathcal{C}$ still exists. If P is then placed at ℓ , yielding σ , then $\mathcal{C}$ still exists unless it had enclosed exactly one unoccupied location, ℓ . However, this is not possible as any cycle in $\sigma' \setminus P$ encircling ℓ would also necessarily encircle neighboring unoccupied location ℓ' . This implies $\mathcal{C}$ is present in σ , a contradiction, so we conclude any cycle in σ' must contain P .

By the conditions in Step 4 of Markov chain $\mathcal{M}$ that must be met before a move occurs, particle P necessarily has fewer than five neighbors in σ and locations ℓ and ℓ' satisfy Property 1 or Property 2. First, suppose they satisfy Property 2. While P might momentarily create a cycle when it expands to occupy both locations ℓ and ℓ' , it will then contract to location ℓ' . Suppose P is part of some cycle $\mathcal{C}$ in σ' . Before and after P on $\mathcal{C}$ are some neighbors N_1 and N_2 of P . By Property 2, N_1 and N_2 are connected by a path in $N(\ell')$, which doesn't contain P . Replacing path $N_1 - P - N_2$ in cycle $\mathcal{C}$ by this path in $N(\ell')$ yields a (not necessarily simple) cycle $\mathcal{C}'$ in σ' not containing P , a contradiction.

Suppose ℓ and ℓ' satisfy Property 1; recall location ℓ has less than five neighbors in σ . Suppose there exists a cycle $\mathcal{C}$ in σ' , which by definition encircles at least one unoccupied location. If ℓ is unoccupied inside $\mathcal{C}$, then so is at least one of its neighbors; we conclude $\mathcal{C}$ encircles some unoccupied location $\ell'' \neq \ell$. Let N_1 and N_2 be the particles on cycle $\mathcal{C}$ before and after P . If there exists a path between N_1 and N_2 in $N(\ell')$, the argument in the previous paragraph applies, so we suppose this is not the case. It must be, without loss of generality, that $|\mathbb{S}| = 2$ and there exist paths in $N(\ell \cup \ell')$ from N_1 to $S_1 \in \mathbb{S}$ and from N_2 to $S_2 \in \mathbb{S}$, with $S_1 \neq S_2$. There then exists cycle $\mathcal{C}'$ in σ , obtained from $\mathcal{C}$ by replacing path $N_1 - P - N_2$, where P is in location ℓ' , with the path $N_1 - \dots - S_1 - P - S_2 - \dots - N_2$, where P is in location ℓ . This is a valid (not necessarily simple) cycle in σ , as it still encircles unoccupied location $\ell'' \neq \ell$. This is also a contradiction, so in all cases we find σ' has no cycles. $\square$

3.4 Ergodicity of Markov chain $\mathcal{M}$

We next show the carefully-defined moves of $\mathcal{M}$ suffice to move from any configuration σ to any other configuration τ , necessary for showing $\mathcal{M}$ is ergodic and thus has a unique stationary distribution.

We emphasize the details of this proof are far from trivial, and occupy ten pages in the full version of this paper*. Figure 2 illustrates one difficulty. It depicts a particle configuration for which there exist no valid moves satisfying Property 1; the only valid moves satisfy Property 2. Thus if moves satisfying Property 2 are not included, the state space of $\mathcal{M}$ is not connected. Our approach relies critically on moves satisfying Property 2.

Here we outline and explain our approach to proving $\mathcal{M}$ is ergodic. At a high level, we prove for any configuration σ there exists a sequence of valid particle moves transforming σ into a straight line, and then prove $\mathcal{M}$ is reversible, implying for any other configuration τ there exists a sequence of valid particle moves transforming that straight line into τ .




Figure 2: A particle configuration for which all valid moves of Markov chain $\mathcal{M}$ satisfy Property 2; no particle has a valid move satisfying Property 1 (darker lines represent induced edges of the system). This demonstrates the subtlety of the Markov chain rules we have defined.

Given a particle configuration σ with lowest leftmost particle S , we find a sequence of moves transforming σ into a line of particles stretching down and left from S ; particle S never moves. We traverse the perimeter of σ , starting at S , finding particles which we can *eliminate*, or move to this line. We begin with a crucial definition.

DEFINITION 3.3. *An unoccupied location ℓ in Γ is a gap of configuration σ if placing an additional particle P at ℓ results in a hole in configuration $\sigma \cup P$.*

LEMMA 3.4. *There exists a sequence of particle moves transforming any configuration σ into a straight line.*

PROOF. Traverse the perimeter of σ starting from lowest leftmost particle S . When possible, particles encountered move via a sequence of valid particle moves backwards along the perimeter to S , and then into a line stretching down and left from S , eliminating them. This is possible precisely when particles can move along the perimeter to S without ever occupying a gap location. We prove if $\mathcal{W}$ is the longest clockwise walk around the perimeter of σ starting at S that is not adjacent to any gaps, then it is possible to either eliminate one particle or execute a valid sequence of particle moves that lengthens walk $\mathcal{W}$. In fewer than $2n^2$ iterations, all particles are eliminated: a particle must be eliminated at least once every $2n$ steps, as the length of $\mathcal{W}$ can increase at most $p(\sigma) < 2n$ times in a row. After this, all particles form a straight line stretching down and left from S . $\square$

We present one more lemma before proving $\mathcal{M}$ is irreducible. Recall $P(\sigma, \tau)$ is the probability of moving from configuration σ to state τ in one step of $\mathcal{M}$.

LEMMA 3.5. *For any two configurations σ and τ in Ω , if $P(\sigma, \tau) > 0$, then $P(\tau, \sigma) > 0$.*

PROOF. Let $\sigma, \tau \in \Omega$ be any two configurations such that $P(\sigma, \tau) > 0$. This means σ and τ differ by one particle P at location ℓ in σ and at adjacent location ℓ' in τ .

Note in τ , particle P at location ℓ' has at most four neighbors. This is because $\ell \in n(\ell')$ is unoccupied as particle P is instead at ℓ' , and at least one other location in $n(\ell')$ is unoccupied as otherwise ℓ' would have been a hole in σ , impossible by Lemma 3.2. Because $P(\sigma, \tau) > 0$, Property 1 or Property 2 must hold for ℓ and ℓ' . Both properties are symmetric with regard to the role played by ℓ and ℓ' . If Markov chain $\mathcal{M}$, in state τ , selects in Step 1 particle P , location $\ell \in n(P)$, and a sufficiently small probability q , then because Conditions 1, 2, and 3 are necessarily satisfied, particle P moves to location ℓ . This proves $P(\tau, \sigma) > 0$. $\square$

LEMMA 3.6. *Markov chain $\mathcal{M}$ connects the state space of all connected configurations without holes.*

PROOF. Let σ and τ be any two connected configurations of n particles with no holes. By Lemma 3.4, there exists a sequence of moves transforming σ into a line with slope $1/\sqrt{3}$. By Lemmas 3.4 and 3.5, there exists a sequence of valid moves transforming this line into τ . $\square$

COROLLARY 3.7. *$\mathcal{M}$ is ergodic.*

PROOF. By Lemma 3.6, the state space is connected. $\mathcal{M}$ is aperiodic as at each iteration there is a probability of at least $1/6$ that no move is made. Therefore $\mathcal{M}$ is ergodic. $\square$

3.5 The stationary distribution π of $\mathcal{M}$

We now know the Markov chain $\mathcal{M}$ is ergodic and finite, so its stationary distribution is unique.

LEMMA 3.8. *The stationary distribution π of $\mathcal{M}$ is*

$$\pi(\sigma) = \lambda^{t(\sigma)} / Z,$$

where $Z = \sum_{\sigma} \lambda^{t(\sigma)}$ is the normalizing constant.

PROOF. We confirm π is the stationary distribution by detailed balance. Let σ and τ be configurations in Ω with $\sigma \neq \tau$ such that $P(\sigma, \tau) > 0$. By Lemma 3.5, also $P(\tau, \sigma) > 0$. Suppose particle P moves from location ℓ in σ to neighboring location ℓ' in τ . Let t be the number of triangles on which Q is incident when it is in location ℓ , and let t' be that number when P is in location ℓ' . This implies $t(\sigma) - t(\tau) = t - t'$. Without loss of generality, let $t' < t$. We see

$$P(\sigma, \tau) = \frac{1}{n} \cdot \frac{1}{6} \cdot \lambda^{t'-t} \quad \text{and} \quad P(\tau, \sigma) = \frac{1}{n} \cdot \frac{1}{6} \cdot 1.$$

We now show σ and τ satisfy the detailed balance condition:

$$\pi(\sigma)P(\sigma, \tau) = \frac{\lambda^{t(\sigma)}}{Z} \frac{\lambda^{t'-t}}{6n} = \frac{\lambda^{t(\tau)}}{Z \cdot 6n} = \pi(\tau)P(\tau, \sigma).$$

We conclude π is the stationary distribution of $\mathcal{M}$. $\square$

While it is natural to assume maximizing the number of triangles in a particle configuration results in more compression, here we formalize this. We prove π can also be expressed in terms of perimeter, which implies $\mathcal{M}$ converges to a distribution weighted by the perimeter of configurations, a global characteristic, even though the probability of any particle move is determined only by local information.

COROLLARY 3.9. *The stationary distribution π of $\mathcal{M}$ is*

$$\pi(\sigma) = \lambda^{-p(\sigma)} / Z,$$

where $Z = \sum_{\sigma} \lambda^{-p(\sigma)}$ is the normalizing constant.

PROOF. We use Lemma 2.3 and Lemma 3.8:

$$\begin{aligned} \pi(\sigma) &= \frac{\lambda^{t(\sigma)}}{\sum_{\sigma} \lambda^{t(\sigma)}} = \frac{\lambda^{2n-p(\sigma)-2}}{\sum_{\sigma} \lambda^{2n-p(\sigma)-2}} \\ &= \frac{\lambda^{2n-2}}{\lambda^{2n-2}} \cdot \frac{\lambda^{-p(\sigma)}}{\sum_{\sigma} \lambda^{-p(\sigma)}} = \frac{\lambda^{-p(\sigma)}}{\sum_{\sigma} \lambda^{-p(\sigma)}}. \quad \square \end{aligned}$$

The stationary distribution of $\mathcal{M}$ can also be expressed in terms of edges.

COROLLARY 3.10. *The stationary distribution π of $\mathcal{M}$ is*

$$\pi(\sigma) = \lambda^{e(\sigma)} / Z,$$

where $Z = \sum_{\sigma} \lambda^{e(\sigma)}$ is the normalizing constant.

PROOF. This follows from Lemma 2.4 and Corollary 3.9:

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

3.6 Convergence Time of Markov Chain $\mathcal{M}$

We prove in Section 4 that when $\lambda > 5$, if Markov chain $\mathcal{M}$ has converged to its stationary distribution, then with all but exponentially small probability the particle system will be compressed. However, we do not give explicit bounds on the time required for this to occur; we give experimental evidence of convergence times in Section 6, but believe proving rigorous bounds will be challenging.

A common measure of convergence time of a Markov chain is the *mixing time*, the number of iterations until the distribution is within total variation distance ε of the stationary distribution, starting from the worst initial configuration. Bounding the mixing time of a Markov chain achieving compression is likely to be challenging because of the similarity to physical systems such as the Ising and Potts models, common models for ferromagnetism. Algorithms that perform local updates are known to require exponential time for many of these models precisely because of a type of compression of the systems [18]. However, mixing time most likely is not the correct measure of our algorithm's convergence. Even if it takes exponential time for $\mathcal{M}$ to converge to its stationary distribution, which is certainly plausible, it may be true that the particles achieve α -compression after only a polynomial number of steps. In fact, based on simulations, it appears compression occurs in polynomial time; doubling the number of particles consistently results in a ten-fold increase in iterations until compression starting from a straight line of n particles, so we conjecture the number of iterations until compression occurs is close to $O(n^{3.3})$.

4. ACHIEVING COMPRESSION

If $\mathcal{M}$ executes long enough, it will converge to its stationary distribution π ; we will use the expression of π given in Corollary 3.9. To simplify notation, we define the *weight* of a configuration σ to be $w(\sigma) = \lambda^{-p(\sigma)}$. For a set $S \subseteq \Omega$, we define $w(S)$ as the sum of the weights of all configurations in S . We now show that, provided λ and n are large enough, with all but exponentially small probability if $\mathcal{M}$ is at stationarity then the particles are in an α -compressed configuration. Constant $\alpha > 1$ can be as close to 1 as desired, though smaller α requires larger λ . We begin with a crucial counting lemma.

LEMMA 4.1. *The number of connected configurations with no holes and perimeter k is at most 5^k .*

PROOF. Consider any configuration σ . Let $\mathcal{W}$ be a clockwise traversal of its perimeter, beginning at the lowest leftmost particle of σ ; for every edge e traversed in $\mathcal{W}$, the location left of e is unoccupied. At each step of $\mathcal{W}$, the perimeter can continue straight, turn left by 60 degrees, turn right by 60 degrees, turn right by 120 degrees, or turn right by 180 degrees; the perimeter can never turn left by 120 or 180 degrees. Thus, at any point there are at most five possible locations for the next particle on $\mathcal{W}$. We conclude that there are fewer than 5^k configurations of perimeter k , each specified by the directions of the turns on $\mathcal{W}$. $\square$

We note this bound is not tight; see the full version* for stronger bounds, which lead directly to lower values of λ^* below. We now prove our main result.

THEOREM 4.2. *For any $\alpha > 1$, there exists $\lambda^* = 5^{\frac{\alpha}{\alpha-1}}$, $n^* \geq 0$, and $\gamma < 1$ such that for all $\lambda > \lambda^*$ and $n > n^*$, the probability that a random sample σ drawn according to the stationary distribution π of $\mathcal{M}$ is not α -compressed is exponentially small:*

$$\mathbb{P}(p(\sigma) \geq \alpha \cdot p_{\min}) < \gamma^{\sqrt{n}}.$$

PROOF. Let S_α be the set of configurations of perimeter at least $\alpha \cdot p_{\min}$. Let $\sigma_{\min}$ be a configuration of n particles achieving the minimum perimeter $p_{\min}$. We show

$$\pi(S_\alpha) = \frac{w(S_\alpha)}{Z} < \frac{w(S_\alpha)}{w(\sigma_{\min})} \leq \gamma^{\sqrt{n}}.$$

The first equality is the definition of π ; the next inequality follows from the definitions of Z and w . We focus on the last inequality. Using Lemma 4.1 and noting the weight of any σ with $p(\sigma) = k$ is $\lambda^{-p(\sigma)} = \lambda^{-k}$, we sum over S_α :

$$\begin{aligned} \frac{w(S_\alpha)}{w(\sigma_{\min})} &\leq \frac{\sum_{k=\lceil \alpha \cdot p_{\min} \rceil}^{2n-2} 5^k \lambda^{-k}}{\lambda^{-p_{\min}}} \\ &= \sum_{k=\lceil \alpha \cdot p_{\min} \rceil}^{2n-2} 5^{(1-\log_5 \lambda)k + (\log_5 \lambda)p_{\min}}. \end{aligned}$$

Using the inequality $p_{\min} \leq k/\alpha$, it follows that

$$\begin{aligned} \frac{w(S_\alpha)}{w(\sigma_{\min})} &\leq \sum_{k=\lceil \alpha \cdot p_{\min} \rceil}^{2n-2} 5^{(1-\log_5 \lambda)k + (\log_5 \lambda)(k/\alpha)} \\ &= \sum_{k=\lceil \alpha \cdot p_{\min} \rceil}^{2n-2} 5^{(1-(1-1/\alpha)\log_5 \lambda)k}. \end{aligned}$$

As $\lambda > \lambda^* = 5^{\frac{\alpha}{\alpha-1}}$, then $-c_1 := 1 - (1-1/\alpha)\log_5 \lambda < 0$. Also $k \geq \alpha \cdot p_{\min}$, and by Lemma 2.1, $p_{\min} > \sqrt{n}$, so

$$\frac{w(S_\alpha)}{w(\sigma_{\min})} \leq \sum_{k=\lceil \alpha \cdot p_{\min} \rceil}^{2n-2} 5^{-c_1 \cdot 2\alpha \sqrt{n}} \leq 2n (5^{-2\alpha c_1})^{\sqrt{n}}.$$

There exists $\gamma < 1$ and n^* such that for all $n \geq n^*$,

$$\begin{aligned} \mathbb{P}(p(\sigma) \geq \alpha \cdot p_{\min}) &= \pi(S_\alpha) \leq \frac{w(S_\alpha)}{w(\sigma_{\min})} \\ &\leq 2n (5^{-2\alpha c_1})^{\sqrt{n}} < \gamma^{\sqrt{n}}. \quad \square \end{aligned}$$

COROLLARY 4.3. *For any $\lambda > 5$, there exists a value $\alpha = \log_5 \lambda / (\log_5 \lambda - 1)$, $n^* \geq 0$, and $\gamma < 1$ such that for all $n > n^*$, a random sample σ drawn according to the stationary distribution π of $\mathcal{M}$ satisfies*

$$\mathbb{P}(p(\sigma) \geq \alpha \cdot p_{\min}) < \gamma^{\sqrt{n}}.$$

5. USING $\mathcal{M}$ FOR EXPANSION

A nice feature of our algorithm is that it also provably achieves particle expansion for different values of bias parameter λ . We say a configuration σ is β -expanded for some $\beta < 1$ if $p(\sigma) > \beta \cdot p_{\max}$, where $p_{\max} = 2n - 2$. We note as $p_{\max} = \Theta(n)$ and $p_{\min} = \Theta(\sqrt{n})$ for a system of n particles, β -expansion and α -compression for any constants β and α are mutually exclusive for sufficiently large n . We

show that, provided n is large enough, for all $0 < \lambda < \sqrt{2}$ there is a constant β such that with all but exponentially small probability, if $\mathcal{M}$ is at stationarity then the particles are β -expanded. This is notable because it implies, counter-intuitively, that $\lambda > 1$ is not sufficient to guarantee particle compression as one might first guess.

THEOREM 5.1. *For all $0 < \beta < 1$, there exists $\lambda^* = \lambda^*(\beta) < \sqrt{2}$, $n^* \geq 0$, and $\gamma < 1$ such that for all $\lambda < \lambda^*$ and $n > n^*$, the particles achieve β -expansion with all but exponentially small probability: for a configuration σ drawn at random according to stationary distribution π ,*

$$\mathbb{P}(p(\sigma) < \beta \cdot p_{\max}) \leq \gamma^{\sqrt{n}}.$$

PROOF. We let

$$\lambda < \lambda^* = 5^{\frac{\beta - \frac{\log_5 2}{2}}{\beta - 1}} < 5^{\frac{0 - \frac{\log_5 2}{2}}{0 - 1}} = \sqrt{2}.$$

Let S_β be the set of configurations σ with $p(\sigma) \leq \beta \cdot p_{\max}$. Let $\sigma_{\max}$ be a configuration achieving maximum perimeter $p_{\max} = 2n - 2$. Note the number N of configurations achieving this maximum is at least $2^{n-1} = 2^{p_{\max}/2}$; this is the number of paths where every step is up or up-right, which have no triangles and thus maximum perimeter. We show

$$\pi(S_\beta) = \frac{w(S_\beta)}{Z} < \frac{w(S_\beta)}{N \cdot w(\sigma_{\max})} \leq \gamma^{\sqrt{n}}.$$

We focus on proving the last of the inequalities above. Applying Lemma 4.1, we see

$$\begin{aligned} \frac{w(S_\beta)}{N \cdot w(\sigma_{\max})} &\leq \frac{\sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} 5^k \lambda^{-k}}{2^{p_{\max}/2} \cdot \lambda^{-p_{\max}}} \\ &= \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} 5^k \lambda^{-k} \left(\frac{\lambda}{\sqrt{2}} \right)^{p_{\max}}. \end{aligned}$$

Recalling that $\lambda < \sqrt{2}$ and $k < \beta \cdot p_{\max}$, we see

$$\begin{aligned} \frac{w(S_\beta)}{N \cdot w(\sigma_{\max})} &\leq \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} 5^k \lambda^{-k} \left(\frac{\lambda}{\sqrt{2}} \right)^{\frac{k}{\beta}} \\ &= \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} 5^{k(1 - \log_5 \lambda + \frac{\log_5 \lambda}{\beta} - \frac{\log_5 2}{2\beta})}. \end{aligned}$$

Let c_3 be such that

$$-c_3 = 1 - \log_5 \lambda + \frac{\log_5 \lambda}{\beta} - \frac{\log_5 2}{2\beta}.$$

We want $-c_3 < 0$, so we solve for λ and see $-c_3 < 0$ precisely when $\lambda < \lambda^*$, a condition we know to hold. It follows that

$$\begin{aligned} \frac{w(S_\beta)}{N \cdot w(\sigma_{\max})} &\leq \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} 5^{-c_3 k} \\ &\leq \sum_{k=p_{\min}}^{\lfloor \beta \cdot p_{\max} \rfloor} 5^{-c_3 \cdot 2\sqrt{n}} \leq 2n 5^{-2c_3 \sqrt{n}}. \end{aligned}$$

Hence there exists $n^* \geq 0$ and $\gamma < 1$ such that for all $n \geq n^*$,

$$\begin{aligned} \mathbb{P}(p(\sigma) \leq \beta \cdot p_{\max}) &= \pi(S_\beta) \leq \frac{w(S_\beta)}{N \cdot w(\sigma_{\max})} \\ &\leq 2n 5^{-c_3 \cdot 2\sqrt{n}} < \gamma^{\sqrt{n}}. \quad \square \end{aligned}$$




Figure 3: 100 particles in a line with occupied edges drawn, after (a) 1 million, (b) 2 million, (c) 3 million, (d) 4 million, and (e) 5 million iterations of $\mathcal{M}$ with bias $\lambda = 4$.

COROLLARY 5.2. *For all $\lambda < \sqrt{2}$, there exists a constant $0 < \beta < 1$ such that with all but exponentially small probability a sample drawn according to stationary distribution π of $\mathcal{M}$ is β -expanded.*

PROOF. This follows immediately from the previous theorem; for $\lambda < \sqrt{2}$, there exists β such that $\lambda < \lambda^*(\beta)$. $\square$

6. SIMULATIONS

In practice, Markov chain $\mathcal{M}$ yields good compression, even beyond the values of λ for which our proofs apply. We simulated $\mathcal{M}$ for $\lambda = 4$ on 100 particles that began in a line; the configurations after 1, 2, 3, 4, and 5 million steps of $\mathcal{M}$ are shown in Figure 3.

In contrast, $\lambda = 2$, while still favoring particles forming triangles, does not appear to yield compression; see Figure 4, where even after 20 million simulated steps of $\mathcal{M}$, the particles have not compressed. We conjecture there is a phase transition in λ , i.e., a critical value λ_c such that for all $\lambda < \lambda_c$ the particles do not compress and for all $\lambda > \lambda_c$ they do compress. Such phase transitions exist for similar statistical physics models (e.g., [2]). Our proofs indicate if λ_c exists, then $\sqrt{2} \leq \lambda_c \leq 5$; simulations suggest $2 < \lambda_c < 4$.

7. REFERENCES

- [1] E. Bampas, J. Czyzowicz, L. Gąsieniec, D. Ilcinkas, and A. Labourel. Almost optimal asynchronous rendezvous in infinite multidimensional grids. In *Distributed Comp.: 24th Int. Symp., DISC 2010*, pages 297–311, 2010.
- [2] C. Borgs, J.T. Chayes, A. Frieze, J.H. Kim, P. Tetali, E. Vigoda, and V.H. Vu. Torpid mixing of some MCMC algorithms in statistical physics. In *40th IEEE Symp. on Found. of Comp. Sci., FOCS 1999*, pages 218–229, 1999.
- [3] S. Camazine, K.P. Visscher, J. Finley, and S.R. Vetter. House-hunting by honey bee swarms: Collective decisions and individual behaviors. *Insectes Sociaux*, 46:348–360.
- [4] A. Chavoya and Y. Duthen. Using a genetic algorithm to evolve cellular automata for 2D/3D computational development. In *Genetic and Evolut. Comp. Conf., GECCO 2006*.
- [5] Z. Derakhshandeh, R. Gmyr, A.W. Richa, C. Scheideler, and T. Strothmann. An algorithmic framework for shape formation problems in self-organizing particle systems. In *Proc. of the 2nd Ann. Int. Conf. on Nanoscale Computing and Comm., NANOCOM' 15*, pages 21:1–21:2, 2015.
- [6] Z. Derakhshandeh, R. Gmyr, A.W. Richa, C. Scheideler, and T. Strothmann. Universal shape formation for programmable matter. In *To appear, 28th ACM Symp. on Parallelism in Alg. and Arch., SPAA '16*, 2016. To appear.
- [7] Z. Derakhshandeh, R. Gmyr, T. Strothmann, R.A. Bazzi, A.W. Richa, and C. Scheideler. Leader election and shape formation with self-organizing programmable matter. In *21st DNA Comp. and Molec. Prog., DNA 21*, pages 117–132, 2015.
- [8] A. Deutsch and S. Dormann. *Cellular Automaton Modeling of Biological Pattern Formation: Characterization, Applications, and Analysis*. Modeling and Simulation in Science, Eng. and Technology. Birkhäuser Boston, 2007.
- [9] R.L. Dobrushin. The problem of uniqueness of a gibbsian random field and the problem of phase transitions. *Functional Analysis and Its Applications*, 2:302–312, 1968.
- [10] W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, 1968.
- [11] P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer. Arbitrary pattern formation by asynchronous, anonymous, oblivious robots. *Theoretical Computer Science*, 407:412–447, 2008.
- [12] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57:97–109, 1970.
- [13] R. Jeanson, C. Rivault, J.L. Deneubourg, S. Blanco, R. Fournier, C. Jost, and G. Theraulaz. Self-organized aggregation in cockroaches. *Animal Behaviour*, 69:169 – 180, 2005.
- [14] N. Lynch. *Distributed Algorithms*. Morgan Kaufman, 1996.
- [15] N.J. Mlot, C.A. Tovey, and D.L. Hu. Fire ants self-assemble into waterproof rafts to survive floods. *Proc. of the National Academy of Sci.*, 108:7669–7673, 2011.
- [16] C. Rivault and A. Cloarec. Cockroach aggregation: Discrimination between strain odours in *Blattella germanica*. *Animal Behaviour*, 55:177–184, 1998.
- [17] M. Rubenstein, A. Cornejo, and R. Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345:795–799, 2014.
- [18] L.E. Thomas. Bounds on the mass gap for finite volume stochastic Ising models at low temperature. *Comm. Math. Phys.*, 126:1–11, 1989.
- [19] E. Winfree, F. Liu, L.A. Wenzler, and N.C. Seeman. Design and self-assembly of two-dimensional DNA crystals. *Nature*, 394(6693):539–544, 1998.
- [20] D. Woods, H.L. Chen, S. Goodfriend, N. Dabby, E. Winfree, and P. Yin. Active self-assembly of algorithmic shapes and patterns in polylogarithmic time. In *Proc. of the 4th Conf. on Innovations in Theoretical Computer Science*, pages 353–354, 2013.




Figure 4: 100 particles after (a) 10 million and (b) 20 million iterations of $\mathcal{M}$ with bias $\lambda = 2$, starting with all particles in a diamond shape of side lengths 10.