Single Bridge Formation in Self-Organizing Particle Systems

Shunhao Oh[1]
School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Joseph L. Briones
School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ, USA

Jacob Calvert
School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Noah Egan
School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Dana Randall
School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Andréa W. Richa
School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ, USA

Abstract

Local interactions of uncoordinated individuals produce the collective behaviors of many biological systems, inspiring much of the current research in programmable matter. A striking example is the spontaneous assembly of fire ants into "bridges" comprising their own bodies to traverse obstacles and reach sources of food. Experiments and simulations suggest that, remarkably, these ants always form one bridge – instead of multiple, competing bridges – despite a lack of central coordination. We argue that the reliable formation of a single bridge does not require sophistication on behalf of the individuals by provably reproducing this behavior in a self-organizing particle system. We show that the formation of a single bridge by the particles is a statistical inevitability of their preferences to move in a particular direction, such as toward a food source, and their preference for more neighbors. Two parameters, η and β, reflect the strengths of these preferences and determine the Gibbs stationary measure of the corresponding particle system's Markov chain dynamics. We show that a single bridge almost certainly forms when η and β are sufficiently large. Our proof introduces an auxiliary Markov chain, called an "occupancy chain," that captures only the significant, global changes to the system. Through the occupancy chain, we abstract away information about the motion of individual particles, but we gain a more direct means of analyzing their collective behavior. Such abstractions provide a promising new direction for understanding many other systems of programmable matter.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Self-organizing particle systems, programmable matter, bridging, jump chain

Digital Object Identifier 10.4230/LIPIcs.DISC.2024.34

Related Version *arXiv Version*: `https://arxiv.org/abs/2408.10830` [17]

1 Introduction

Ants are known to collaborate to form elaborate structures. Army ants of genus *Eiticon* self-assemble and disperse while foraging to form shortcuts in order to more efficiently reach food [19]. Weaver ants of species *Oecophylla longinoda* bind leaves together when building

[1] Corresponding author.

their nests and can form long chains [13]. In this work, we take inspiration from the behavior of fire ants of species *Solenopsis invicta*, which can self-assemble into floating rafts and structural bridges [16]. These bridges are structures composed of the ants as they entangle themselves and allow other ants a medium over which they can reach a target, such as food. It is not well understood why individual ants sacrifice themselves by becoming entangled with others in unfavorable circumstances such as suspension over water, but such occurrences are regularly accomplished without any centralized coordination.

Recent laboratory experiments at the University of Georgia [24] were the first to replicate the more general fire ant bridging found in nature by placing food in the center of a bowl filled with water to see how hungry ants will coordinate. Initially, ants surround the bowl looking for food, but soon they begin to explore the water, begin building multiple structures extending from the edge of the bowl. Eventually some structures start to reach the food, but over time only a *single bridge* persists, with the other ants traversing the bridge to bring food back to the nest. Biologists have questioned what coordination allows the ants to always form a single bridge, streamlining their initial structures to minimize the number of ants sacrificed for the bridge while enabling the remaining ants to productively harvest food for the nest [19]. Physicists, biologists and computer scientists [4] have developed an agent-based model that mimics this behavior with locally interacting particles that primarily depends on two parameters: one capturing their individual awareness of the food source, based on proximity, and another that captures the rigidity of the bridge structure by measuring an interparticle attraction among all nearest neighbor pairs, which has been shown to suffice to encourage collectives to aggregate [4]. This physical model also allows ants to climb over one another and incorporates additional physical parameters such as the meniscus effects at the edge of the bowl and effects of small clusters disconnecting. Their simulations reliably replicate the tendency for exactly one bridge to form, but lack an explanation because the biological and simulated protocols are too difficult to analyze in order to connect the local behavior of individuals with the complex coordinated structure formed by the ensemble.

We show that the emergence of at most one bridge, as observed in ant experiments, follows from a statistical inevitability, and not from the sophistication of the individuals. Here, rather than tracking the movements of individual particles (or ants) over time according to a local Markov chain, we define a related auxiliary *occupancy chain* that captures only the changes to which sites are occupied. The states of the occupancy chain record the sites where one or more particles currently exist and the transitions capture the stochastic changes to the occupied sites over time, to determine long-term behavior. This approach of using occupancy chains allow us to apply techniques from statistical physics to infer properties of contours, and the simplified analysis is likely to be useful in other contexts, such as various fixed magnetization spin systems [8], biased growth processes [9], and collectives arising in swarm robotics responding to directed external stimuli [12].

1.1 Related work

Ant bridging can be viewed as an example of programmable matter, in which researchers seek to create materials that can be programmed to change their physical properties in response to user input and environmental stimuli. Diverse examples of programmable matter – like DNA tiles [22], synthetic cells [11], and reconfigurable modular robots [5] – are linked through their common use of local interactions to produce global functionality beyond the capabilities of individuals. Examples of this phenomenon abound in nature, including honey bees that select nest sites using decentralized recruitment [2], cockroach larvae that aggregate using short-range pheromones [10], and fire ants that assemble their bodies into rafts to

Figure 1 Photos of fire ants, *Solenopsis*, self-assembling over time to form a bridge to reach food placed in the center of a bowl filled with water. Experiments performed by and photo credit to Takao Sasaki and Horace Zeng at the University of Georgia.

survive floods [16]. For this reason, programmable matter is often designed by analogy with biological collectives that exhibit an emergent behavior of interest [20, 18]. However, this form of biomimicry is generally unamenable to rigorously proving that an algorithm executed by the constituents of matter reliably leads to a desired behavior. This analysis is critical to understanding the extent to which algorithms for programmable matter are robust when they are actually deployed [12].

An abstraction of programmable matter known as a *self-organizing particle system* (SOPS) addresses this gap. In the standard amoebot model for SOPS [7, 6], particles exist on the sites (or nodes) of a lattice, with at most one particle per site, and move between sites along lattice edges. Each particle is anonymous (unlabeled), interacts only with particles on adjacent lattice sites, has limited (normally assumed to be constant-size) memory, and does not have access to any global information such as a common direction on the lattice, a coordinate system or the total number of particles. We note that in this paper, we are studying a particle system with possibly multiple particles (ants) per site, but we reduce to an occupancy chain which recovers the "SOPS perspective" by just indicating whether each site is occupied or not. Recent work on SOPS has led to rigorous stochastic distributed algorithms for various tasks by exploring phase changes that lead to desirable collective behavior with high probability. Examples include compression and expansion [4], where particles cluster closely or expand; separation [3], where colored particles would like to separate into clusters of same color particles; aggregation [12], where the system would like to compress but no longer needs to remain connected; shortcut bridging [1]; locomotion [21, 23]; and transport [12]. However, while these algorithms based on local Markov chains provide insight into emergent collective behavior, the proofs guaranteeing the long-term behaviors become prohibitively challenging as the collectives take on increasingly complicated tasks with more sophisticated interactions. For example, while the work in [1] also looked at the bridging behavior of fire ants [19], it did not focus on the bridge formation process. Instead, the paper assumed that a single bridge already existed, anchored at two fixed points (corresponding to the nest and the food), and considered the (simpler) problem of optimizing the placement of the bridge with respect to parameters of relevance.

1.2 Overview of our results

To gain insight into structures that emerge from ant bridging, we consider a discrete model on a finite $w \times h$ domain $\Lambda = \Lambda(w, h)$ of the triangular lattice, with periodic boundary conditions in the horizontal direction, representing the circular bowl containing food. It

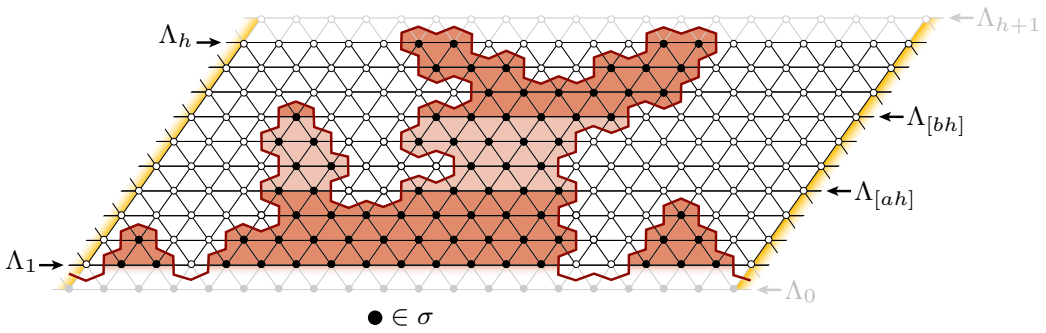

Figure 2 A domain Λ of the triangular lattice of width of $w = 20$ and height of $h = 10$ with periodic boundary conditions identifying the ends of rows (indicated by yellow shading). The configuration σ (black dots) has multiple (a, b)-bridges for $(a, b) = (0.4, 0.7)$, but not for $(a, b) = (0.4, 0.8)$.

will be convenient to think of Λ as a cylindrical domain comprised of h layers (or cycles), corresponding to the h rows Λ_1 through Λ_h, each of width w (see Figure 2). We model the motion of a finite set of n particles entering Λ at Λ_1, and growing a connected component by occupying sites on the interior of Λ in order to eventually reach the food, which we assume can be found at every site on Λ_h. Particles may move onto empty sites or walk upon sites currently occupied by other particles – we intentionally abstract away the specific dynamics by directly analyzing the occupancy chain, which only tracks changes to the set σ of occupied sites in Λ. In analogy with fire ant bridging, a step of the occupancy chain can be interpreted as a lone ant on the boundary of σ climbing over a neighboring ant, thus creating an unoccupied site, or as an ant concluding its walk over other ants to occupy a new site adjacent to the previous boundary. For simplicity, we assume that the aspect ratio $\alpha = w/h$ is a constant and that n scales linearly with the area of the region, writing $n = \lfloor \rho h^2 \rfloor$ for some constant $\rho > 0$, as these are the conditions under which bridges may form.

The occupancy chain captures the relative preferences of particles to move toward Λ_h (e.g., due to the scent of the food source placed along Λ_h), parameterized by $\eta > 0$, and their preference for locations with more neighbors, increasing the rigidity of the connected structure. The balance of these preferences is expressed through a function on configurations called the *Hamiltonian*, which is related to the stationary probability of configurations and expresses the overall "fitness" of a configuration, with each particle contributing its piece.

For a configuration σ and a site $v \in \sigma$, let $S(v)$ be the *intensity of the scent* at v, which is any nonincreasing function of graph distance $d(v)$ from v to Λ_h. Let $B_\sigma(v)$ be the *number unoccupied sites that neighbor* v. Then the Hamiltonian H of a configuration σ is

$$H(\sigma) = -\sum_{v \in \sigma} \left(\eta S(v) - B_\sigma(v)\right), \tag{1}$$

where $\eta > 0$ is a parameter describing how much a particle favors a stronger scent relative to how much it desires more neighbors. A second parameter β, the interparticle attraction, captures the strength of the Hamiltonian. These induce a probability distribution π over the set of valid configurations Ω (defined in Section 2.1) known as the Gibbs distribution, where

$$\pi(\sigma) = \exp\left(-\beta H(\sigma)\right)/Z, \quad \text{for } \sigma \in \Omega \tag{2}$$

in terms of the normalization constant $Z = \sum_{\tau \in \Omega} e^{-\beta H(\tau)}$. Note we have taken some liberties in defining the occupancy chain by assuming that each particle has been walking on the current bridge structure sufficiently long that it is near equilibrium.

We may now talk about the shape of the contour between the set of occupied and unoccupied sites on Λ in order to determine the number of bridges at equilibrium. We introduce a weak definition of "multiple bridges" to show that any contours that unnecessarily "backtrack" a nontrivial amount from *any* layer will be extremely unlikely when the interparticle attraction β is large enough. We say that a configuration σ has *multiple* (a, b)*-bridges* for real numbers $0 < a < b < 1$ if the restriction of σ to layers $\lfloor ah \rfloor$ and greater contains at least two components (on the lattice Λ) that reach layer $\lfloor bh \rfloor$. This is formally stated later in Definition 4 . Intuitively, one can imagine such configurations as ones that extend out to layer $\Lambda_{\lfloor bh \rfloor}$, retract back to $\Lambda_{\lfloor ah \rfloor}$, then extend out to $\Lambda_{\lfloor bh \rfloor}$ a second time.

Our first main theorem states that, on any $w \times h$ region Λ with constant aspect ration α, regardless of the scent parameter η, if the interparticle attraction β is large enough, then for any constant $\epsilon \in (0, 1)$, the *probability of seeing multiple* (a, b)*-bridges where* $b - a \geq \epsilon$ *is always exponentially small in* β *and* h (see Theorem 5 in Section 3.2). Taking $b = 1$ shows that we are unlikely to see a contour that reaches the food multiple times while backtracking a distance ϵ between the intervals in which they touch.

Our second main result concerns whether or not any bridge will be likely to form. This question is much more delicate and depends on the scent parameter η, some basic assumptions about the scent gradient (see Definition 9), and some bounds on the number of particles in the system. If the *scent parameter* η *is too small*, then particles will compress around the boundary and are unlikely to reach the food, so *no bridge will form.* If the *scent parameter is sufficiently large*, then a *single bridge will emerge, but never more.* These conditions are properly outlined in Section 3.2 and stated as Theorem 10 and Corollary 11.

The proofs are motivated by a technique from statistical physics known as a Peierls argument that constructs a mapping from low-probability configurations to high-probability ones, to infer global properties of the distribution π. The key challenge in constructing suitable mappings for our model is the dependence of the scent component S of the Hamiltonian on the distance of sites from Λ_h. As a consequence, standard physics approaches based on simple mappings that rearrange sites is insufficient. We overcome this challenge by introducing the *layer sequence* of a configuration, which counts the number of the configuration's elements in each layer. Its virtue is that it is far easier to describe and analyze transformations of layer sequences than those of configurations, and the layer sequence retains the information necessary to calculate the scent component of the configuration's Hamiltonian. Moreover, while it discards the geometric information that is needed to calculate the configuration's boundary length, we are nevertheless able to obtain useful bounds on the boundary length.

Our choice to model the motion of the contour directly via an occupancy chain, instead of inferring it from dynamics prescribed to individual particles (ants), is a significant and valuable departure from the norm in distributed computing. Without this abstraction tracking only the salient changes to the profile of the ensemble, it would be far more difficult to rigorously analyze the bridge formation behavior. Notably, this abstraction comes at the cost of obscuring the underlying particles' computational capabilities and the local distributed algorithm that govern the particles' dynamics, which we address and clarify in Section 2.2.

The success of SOPS in designing simple, local algorithms that produce complex, global behaviors suggests that the apparently intelligent behaviors of natural collectives may in fact be the inevitable consequences of unremarkable circumstances. We see a broader opportunity here, not only to explain how the intelligent behaviors of natural collectives arise, but also to demonstrate that such an explanation need not assume great intelligence or detailed knowledge on behalf of the individuals. We believe incorporating occupancy chains may be a way to provide insight into other complex collectives.

2 The Bridging Occupancy Chain

In the preceding section, we introduced the domain Λ, the Hamiltonian H, and the probability distribution π that H induces, to the extent necessary to informally state our results. In particular, we mentioned that π is supported on a subset of valid configurations Ω, but deferred a definition. We now introduce an extension of Λ that we use to define Ω, and then we elaborate the terms of the Hamiltonian, before defining π and the occupancy chain.

2.1 The model

Some subsets of the vertices of Λ are not suitable abstractions of the bridges in Figure 1, because they do not extend from Λ_1 or because they contain holes. To make this precise, we extend the domain Λ with two further layers, Λ_0 and Λ_{h+1}, the sites of which we respectively treat as always occupied and always unoccupied (Figure 2). We denote by $\overline{\Lambda}$ this extended domain, which is isomorphic to Λ with a height of $h+2$, as shown in Figure 2. We define the set of valid configurations to be

$$\Omega = \left\{\sigma \subseteq V(\Lambda) : \ |\sigma| \leq n \ \text{ and } \ \sigma \cup V(\Lambda_0) \text{ is simply connected in } \overline{\Lambda}\right\},$$

where $|\sigma|$ denotes the number of elements of σ.

Recall that the Hamiltonian (Equation 1) of a configuration $\sigma \subseteq V(\sigma)$ involves a sum over sites $v \in \sigma$, with contributions from the scent intensity $S(v)$ at v and the number $B_\sigma(v)$ of neighbors of v that are not in σ. To be precise, we define $B_\sigma(v)$ according to

$$B_\sigma(v) = \left|\{u \in V(\Lambda) \setminus \sigma : (u,v) \in E(\Lambda)\}\right|,$$

which serves to penalize configurations that are less compact. Note that, if $\sigma \in \Omega$, then $B_\sigma(v) < 6$ for every $v \in \sigma$. The scent intensity of a site $v \in V(\Lambda)$ is a function that depends only on its distance to the food, which we will assume is stronger the closer a site is to the food. For the sake of the proofs however, we will phrase these as distances $d(v, \Lambda_0)$ from the bottom of the region $\overline{\Lambda}$, noting trivially that $d(v, V(\Lambda_0)) = h - d(v, V(\Lambda_h))$. For convenience of notation, we define a nondecreasing function $S_h : [h] \mapsto \mathbb{R}_{\geq 0}$, and write the scent intensity at a vertex v as

$$S(v) = S_h\big(d(v, V(\Lambda_0))\big).$$

While the function S_h being nondecreasing is sufficient to show in Section 3.1 that multiple bridges are exponentially unlikely, when we discuss whether or not a bridge to the food will form in Section 3.2, we use a slightly narrower class of functions that we call *nondecelerating* (Definition 9). This covers a broad class of functions, including but not limited to:

$$S_h(y) = C_h y^k - D_h \quad \text{or} \quad S_h(y) = C'_h (h-y)^{-k'} - D'_h,$$

where $k, k' \geq 1$, where the values C_h, C'_h normalize the sum of S_h over a "column" of Λ to a constant, and D_h, D'_h are chosen to make $S_h(1) = 0$. The purpose of the normalization is to make the contributions of $S(v)$ and $B_\sigma(v)$ to the Hamiltonian comparable, so η controls whether a bridge forms. In the case of $k, k' = 1$, the first option gives us a linear scent gradient, while the second option represents a scent intensity proportional to the reciprocal of the distance from the food. Section 3.2 discusses the cases when this sum is not normalized.

We define the occupancy chain in the next section and show that it constitutes an aperiodic and irreducible Markov chain on Ω, hence it has a unique stationary distribution on Ω, as given in Equation 2. Observe that by writing $\lambda = e^{\beta} > 0$ and $\gamma = e^{\beta\eta} > 0$, this formulation can be equivalently written as

$$\pi(\sigma) \propto \prod_{v \in \sigma} \left(\lambda^{-B_\sigma(v)} \gamma^{S(v)} \right), \tag{3}$$

where the probability of each configuration can be written as a product of weights of its individual particles, normalized.

Although the Hamiltonian is partly inspired by the Ising model where the preference for having more neighbors can be thought of as the strength of a ferromagnetic interaction between particles, the scent component of the Hamiltonian introduces a directional bias that deviates from classical models, making the analysis considerably more challenging (see, e.g., [9]). Viewing the dynamics through the occupancy chain gives us a succinct tool to derive emergent behaviors of the collective not as easily provable for related agent-based models.

2.2 The algorithm

The bridging occupancy chain algorithm can be briefly described as follows: Initially, the sites of Λ are unoccupied. Given a configuration σ, we select a uniformly random site v of Λ. If v is occupied, then we attempt to remove it from σ. Otherwise, if v is unoccupied and $|\sigma|$ is less than n, then we attempt to add v to σ. The move is rejected if it produces a configuration that does not belong to Ω, i.e., if it is not simply connected in $\overline{\Lambda}$.

To enforce that simple connectivity is maintained at all times, starting from a simply connected configuration, we follow what we call *local simple connectivity* checks, that can be shown to ensure that the configuration as a whole always remains simply connected in $\overline{\Lambda}$ (Lemma 2). Note that the checks below only depend on a site v and its immediately adjacent sites in $\overline{\Lambda}$, which will be important when we discuss the underlying particle dynamics at the end of this section.

▶ **Definition 1** (Local Simple Connectivity). *For a site $v \in V(\Lambda)$, our definition of local simple connectivity depends on whether v is occupied or not. Let $\overline{\mathcal{N}}(\sigma, v)$ denote the extended neighborhood of v in $\sigma \cup V(\Lambda_0)$, where v is any site of Λ.*

1. *If $v \in \sigma$, we say that v is locally simply connected iff $|\overline{\mathcal{N}}(\sigma, v)| \leq 5$ and the induced subgraph $\overline{\Lambda}[\overline{\mathcal{N}}(\sigma, v)]$ is connected.*
2. *If $v \notin \sigma$, we say that v is locally simply connected iff $|\overline{\mathcal{N}}(\sigma, v)| \geq 1$ and the induced graph $\overline{\Lambda}[\overline{\mathcal{N}}(\sigma, v)]$ is connected.*

The intuition behind this definition is that in Case (1) it is safe to remove a locally simply connected site v from σ, since making v unoccupied cannot disconnect its neighborhood in σ (nor create a hole at v); in Case (2), making the site v occupied cannot connect previously disconnected neighboring sites in σ and therefore cannot create a hole. Note that, for every locally simply connected move, the reverse is also a valid move.

Algorithm 1 describes the moves of the occupancy chain. Note that we are using a variant of the Metropolis-Hastings algorithm [15], but it is easy to check that for any two configurations whose symmetric difference is a single vertex, detailed balance will be satisfied.

The following two lemmas prove the correctness of the Bridging Occupancy Chain. We start by proving that the algorithm will keep the configurations simply connected and then show the Bridging Occupancy Chain is irreducible and aperiodic, and therefore converges to the indicated stationary distribution π.

Algorithm 1 Bridging Occupancy Chain.

Let σ denote the current configuration.
$v \leftarrow$ uniformly random site in Λ
$p \leftarrow$ uniformly random number in $[0, 1]$
if v is locally simply connected **then**
  if $v \in \sigma$ and $p \leq \min\{1, \exp(\beta(2B_\sigma(v) - \eta S(v)))\}$ **then**
    Make v unoccupied.
  else if $v \notin \sigma$, $|\sigma| < n$, and $p \leq \min\{1, \exp(-\beta(2B_\sigma(v) - \eta S(v)))\}$ **then**
    Make v occupied.

▶ **Lemma 2** (Maintaining Simple Connectivity). *Adding or removing a locally simply connected site from a configuration $\sigma \in \Omega$ maintains simple connectivity.*

Proof. Let σ be the current configuration, which we define as simply connected if $\sigma \cup V(\Lambda_0)$ is simply connected in $\overline{\Lambda}$. As Λ_0 is treated as always filled, $\sigma \cup V(\Lambda_0)$ is never empty.

The only way the removal of an occupied site $v \in \sigma$ can make the resulting configuration not simply connected is (i) if it introduces a hole, but this would imply that v had six occupied neighbors in $\sigma \cup V(\Lambda_0)$, which is not allowed; or (ii) if it disconnects the configuration, but that would imply that the neighboring occupied sites to v in $\sigma \cup V(\Lambda_0)$ were disconnected, which we also do not allow. On the other hand, the only way the addition of an occupied site v to σ can violate simple connectivity is if it introduces a cycle that was not present in $\sigma \cup V(\Lambda_0)$. This can only happen if the neighbors of v in $\sigma \cup V(\Lambda_0)$ were either the empty set or if they induced a disconnected graph in $\sigma \cup V(\Lambda_0)$, which we again do not allow. The addition of an occupied site to σ cannot disconnect the resulting configuration. ◀

▶ **Lemma 3** (Irreducibility and Aperiodicity). *The Bridging Occupancy Chain is aperiodic and irreducible. Thus it converges to the stationary distribution π given by Equation 2.*

Proof. We first observe that for every locally simply connected move, the reverse is also valid and locally simply connected. Next, we establish irreducibility of the occupancy chain by showing that from every configuration $\sigma \in \Omega$, we can always identify a site $v \in \sigma$ such that making v unoccupied is a valid move. This would imply that we can always reach the empty configuration via valid moves, and therefore we can connect any pair of configurations $\sigma, \tau \in \Omega$ by first making all the sites in σ unoccupied and then making all the sites in τ occupied by following the reversal of the path from τ to the empty configuration.

Let v be any site in σ that has greatest distance from a site in Λ_0 following edges in $\overline{\Lambda}$ with both endpoints occupied. We will argue that v is locally simply connected. Removing v cannot form a hole because v has greatest distance from Λ_0 and therefore cannot have six occupied neighbors. Likewise, making v unoccupied cannot disconnect σ or there would be occupied sites whose shortest path to $V(\Lambda_0)$ passes through v, contradicting v having greatest shortest distance to $V(\Lambda_0)$. Therefore v is locally simply connected.

Finally, we note that when σ is nearly empty, most choices of v will result in self-loops in the chain, so the chain is aperiodic. Hence, since the chain is ergodic (aperiodic and irreducible), it must have a unique stationary distribution, which is given by Equation 2 due to the Metropolis–Hastings transition probabilities satisfying detailed balance. ◀

Now that we have established the occupancy chain dynamics, it remains to discuss the underlying SOPS dynamics that would govern the behavior of the chain. While the occupancy chain abstraction comes at the cost of obscuring the details of the underlying algorithm that

each particle runs, there are indeed local distributed algorithms running independently on constant-memory particles, which we here assume to operate under a sequential scheduler (i.e., where at most one particle is active at any point in time, as in e.g., [4]), that can closely approximate the occupancy chain dynamics. We first note that each iteration of the Bridging Occupancy Chain, and in particular the checks for local simple connectivity within it, can be executed by constant-sized memory agents with only local communication, since the algorithm only depends on the the states of the agents that occupy site v and its immediate neighborhood in Λ.

In the physics agent-based model for the fire-ant bridging biological experiments in [24] that were the original motivation of this work, each particle (ant) performs an independent random walk, biased by the scent gradient as it walks over other particles on the bridge, and biased by both the scent gradient and number of nearest neighbors as it leaves or joins the boundary. The transition probabilities of the occupancy chain are the probabilities with which the bridge structure gains or loses one particle when this particle's walk on the current bridge structure is at stationarity. Moreover, while fine differences in the specific scent gradient affect the motion of the contour, this merely requires that individuals sense, and are acted upon by, the local gradient and does not require them to store and compute with correspondingly fine numbers. Thus, the occupancy chain model is rich enough to capture the long-term behavior of distributed collectives, and can also describe other underlying dynamics because it abstracts details that are irrelevant to our analysis.

3 Equilibrium properties of the occupancy chain

We are now prepared to prove that at most one bridge will form by analyzing the stationary distribution of the occupancy chain. First, we argue that multiple bridges are unlikely if the affinity parameter is sufficiently large. Then, we identify a phase change in the formation of bridges as the scent parameter varies.

3.1 Precluding the formation of multiple bridges

In Section 1.2 we briefly described a configuration with multiple (a, b)-bridges as one that extends out to layer $\Lambda_{\lfloor bh \rfloor}$, retracts back to $\Lambda_{\lfloor ah \rfloor}$, then extends out to $\Lambda_{\lfloor bh \rfloor}$ a second time. We now formally define this, as well as the *depth* of the backtrack for such configurations.

▶ **Definition 4** (Multiple (a, b)-bridges). *For real values $0 < a < b < 1$, a configuration $\sigma \subseteq V(\Lambda)$ has* multiple (a, b)-bridges *if the subgraph of Λ induced by $\bigcup_{k=\lfloor ah \rfloor}^{h} (\sigma \cap V(\Lambda_k))$ contains at least two components that each contain some vertex from $\Lambda_{\lfloor bh \rfloor}$. We call $\epsilon = b - a > 0$ the* depth *of the backtracking.*

To describe the rarity of multiple bridges forming, we consider the event that the configuration has multiple (a, b)-bridges for some pair (a, b) with depth at least $\epsilon > 0$, defined as:

$$\mathsf{MB}_\epsilon = \cup_{a \in (0, 1-\epsilon)} \{\sigma \in \Omega : \sigma \text{ has multiple } (a, a+\epsilon)\text{-bridges}\}. \tag{4}$$

We note that this event includes intervals of length greater than ϵ as a configuration with multiple $(a, a+\epsilon')$-bridges for $\epsilon' > \epsilon$ also has multiple $(a, a+\epsilon)$-bridges. Our next result tells us that the probability of MB_ϵ decays exponentially in β and h, so long as β is large enough in terms of the aspect ratio α and depth ϵ.

▶ **Theorem 5** (No Multiple Bridges). *For every $\epsilon > 0$ and every $w \times h$ region Λ with aspect ratio $\alpha = w/h$, there exists a positive number $\beta_0 = \beta_0(\alpha, \epsilon)$ such that, if $\beta > \beta_0$, then the probability $\pi(\mathsf{MB}_\epsilon)$ of multiple bridges with depth ϵ is at most $e^{-\beta h \epsilon/2}$.*

We emphasize that the quantity β_0 in Theorem 5 does not depend on the constant ρ (recall that $n = \lfloor \rho h^2 \rfloor$) and scent parameter η, even though these parameters affect the stationary distribution. Instead, as the proof explains, MB_ϵ is rare because its occurrence requires that the configuration has a relatively long boundary. While the number of configurations in MB_ϵ grows exponentially in the perimeter of the domain Λ, this number is insufficient to compensate for the configurations' low weight when the boundary length penalty is high. Accordingly, for MB_ϵ to be rare, it suffices for β to be large in terms of α, which controls the perimeter, and ϵ, which controls the boundary length when MB_ϵ occurs.

We now state three lemmas that we will use to prove Theorem 5. The first two of these results are due to the technical machinery of layer sequences, which we mentioned in Section 1.2 and which we formally introduce in Section 3.3. We note that, although we their proofs appear later, they do not rely on any earlier results. We state them in terms of the total boundary length and scent components of the Hamiltonian, which we define for a configuration $\sigma \in \Omega$ as

$$B(\sigma) = \sum_{v \in \sigma} B_\sigma(v) \quad \text{and} \quad S(\sigma) = \sum_{v \in \sigma} S(v).$$

▶ **Lemma 6.** *For every $\sigma \in \Omega$, there exists $\tau \in \Omega$ such that $B(\tau) \leq \min\{B(\sigma), 6w + 4h\}$ and $S(\tau) \geq S(\sigma)$.*

Lemma 6 states that in the case where σ has boundary length greater than $6w+4h$, we can always find another configuration τ that has boundary length at most $6w + 4h$ with at least as great a scent component (we can set $\tau = \sigma$ otherwise). This implies that $H(\sigma) \geq H(\tau)$ for all positive values of β and η. Moreover, the boundary length of τ is bounded above in terms of $w + h$, instead of $w \times h$. The proof of Lemma 6 appears in Section 3.3.

If σ has multiple bridges, then we can further guarantee that there is a configuration τ with a boundary that is strictly shorter by an amount proportional to the depth of the bridges. This is the content of Lemma 7, which we prove in Appendix A.

▶ **Lemma 7.** *For every $\epsilon > 0$ and $\sigma \in \mathsf{MB}_\epsilon$, there exists $\tau \in \Omega$ such that $B(\sigma) \geq B(\tau) + \epsilon h$.*

Together, Lemmas 6 and 7 imply that a configuration σ with multiple bridges can be mapped to a configuration τ that has far greater probability under π because $H(\sigma) - H(\tau) \geq \beta\epsilon h$. This observation is one part of the proof of Theorem 5. The other part is a bound on the number of configurations with a given boundary length.

▶ **Lemma 8.** *The number of configurations with a boundary length of ℓ is at most $2^{\ell+2w}$.*

Proof. Each configuration $\sigma \in \Omega$ of boundary length ℓ is fully described by a collection of self-avoiding paths of total length ℓ on the dual lattice. Each of these paths must start and end on either the top or bottom of Λ, giving $2w$ possible starting points. There are at most 2^{2w} ways to label each starting point as used or unused, and no more than 2^ℓ ways to draw these paths given known starting points as there are at most two directions on the dual lattice to take each step of the path, giving an upper bound of $2^{\ell+2w}$. ◀

With these three lemmas, we are now prepared to prove our first main theorem:

Proof of Theorem 5. We aim to upper bound the probability that the event MB_ϵ occurs, using the fact that we can map the configurations it comprises to configurations with shorter boundaries. More precisely, by Lemma 7, for every $\sigma \in \mathsf{MB}_\epsilon$, there exists $\tau_\sigma \in \Omega$ such that

$$B(\sigma) - B(\tau_\sigma) \geq \epsilon h. \tag{5}$$

In fact, we can assume that τ_σ further satisfies

$$S(\sigma) \leq S(\tau_\sigma) \quad \text{and} \quad B(\tau_\sigma) \leq 6w + 4h \tag{6}$$

because, if it did not, then we could use Lemma 6 to map τ_σ to a configuration that satisfies all three bounds in its place. We partition MB_ϵ according to the difference in (5):

$$\mathsf{MB}_\epsilon = \cup_{k \geq \epsilon h} \mathsf{MB}_{\epsilon,k}, \quad \text{where} \quad \mathsf{MB}_{\epsilon,k} = \{\sigma \in \mathsf{MB}_\epsilon : B(\sigma) - B(\tau_\sigma) = k\}. \tag{7}$$

The probability of each $\sigma \in \mathsf{MB}_{\epsilon,k}$ is exponentially small in k, due to the first bound in (6):

$$\pi(\sigma) \leq \frac{\pi(\sigma)}{\pi(\tau_\sigma)} = e^{-\beta(H(\sigma)-H(\tau_\sigma))} \leq e^{-\beta(B(\sigma)-B(\tau_\sigma))} = e^{-\beta k}. \tag{8}$$

This bound is useful because $|\mathsf{MB}_{\epsilon,k}|$ is at most exponentially large in k. Indeed, due to the second bound in (6), every $\sigma \in \mathsf{MB}_{\epsilon,k}$ has boundary length of at most

$$B(\sigma) = B(\tau_\sigma) + k \leq 6w + 4h + k.$$

There are at most $2^{2w+\ell}$ configurations with a boundary length of exactly ℓ (Lemma 8), hence the number of configurations in $\mathsf{MB}_{\epsilon,k}$ is at most

$$|\mathsf{MB}_{\epsilon,k}| \leq |\{\sigma \in \Omega : B(\sigma) \leq 6w + 4h + k\}| \leq \sum_{\ell \leq 6w+4h+k} 2^{2w+\ell} \leq 4^{7w+4h+k}. \tag{9}$$

We combine the partition of MB_ϵ with the estimates of $\pi(\sigma)$ and $|\mathsf{MB}_{\epsilon,k}|$ to conclude that

$$\pi(\mathsf{MB}_\epsilon) \overset{(7)}{=} \sum_{k \geq \epsilon h} \sum_{\sigma \in \mathsf{MB}_{\epsilon,k}} \pi(\sigma) \overset{(8)}{\leq} \sum_{k \geq \epsilon h} |\mathsf{MB}_{\epsilon,k}| e^{-\beta k} \overset{(9)}{\leq} \sum_{k \geq \epsilon h} 4^{7w+4h+k} e^{-\beta k}.$$

Since $w = \alpha h$, this bound is at most $e^{-\beta h \epsilon/2}$ when β is large enough in terms of α and ϵ. ◀

3.2 Conditions for bridge formation

In this section, we show that whether a bridge to the food forms depends on the balance between the scent and boundary length components of the Hamiltonian. Recall that η controls each particle's desire to move toward the food relative to its desire for more neighbors.

The competition between scent and boundary length is richest when there are not too few or too many particles in the system. When there are too few particles, no bridges can form; when there are too many particles, "bridges" can form without incurring a high boundary cost, due to the periodic boundary conditions. Accordingly, we assume that the number of particles n is $\lfloor \rho h^2 \rfloor$, for a constant ρ that satisfies $1/2 < \rho < \alpha - 2$, in terms of the aspect ratio $\alpha = w/h$ of the domain Λ. Implicitly, we require that α is greater than 2.5.

We write the scent intensity as a function $S_h : [h] \to \mathbb{R}_{\geq 0}$, where $S_h(k)$ represents the intensity of any particle on layer k. Previously, we only assumed that S_h was a nondecreasing function (i.e., nonincreasing in the distance from the food), and showed that as long as β is sufficiently high, at most one bridge forms. Now we assume additional light conditions on S_h, which we believe will apply for most "natural" definitions of the scent intensity function. We use the term *nondecelerating functions* to refer to this broad class of functions.

▶ **Definition 9.** *For a function $S_h : [h] \to \mathbb{R}_{\geq 0}$, let $\Delta S_h(k) = S_h(k+1) - S_h(k)$ denote the discrete derivative of S_h at $k \in [h-1]$. We say that S_h is* nondecelerating *if $S_h(1) = 0$, $\Delta S_h(k) \geq 0$ for all $k \in [h-1]$, and $\Delta(\Delta S_h)(k) \geq 0$ for all $k \in [h-2]$.*

Assuming that the scent function is nondecelerating, our second main result identifies a phase change in the probability that no bridge reaches the food. We define this event as

$$\mathsf{NB} = \{\sigma \in \Omega : \sigma \cap V(\Lambda_h) = \emptyset\}.$$

The phases correspond to ranges of β and η that depend on the *column scent* $\varphi = \sum_{k=1}^{h} S_h(k)$, which represents the sum of the scent intensities over a single "column" of Λ. We consider the case when $\varphi > 0$ is a constant with respect to h, the aspect ratio satisfies $\alpha > 2.5$, and the constant ρ belongs to $(0.5, \alpha - 2)$. We define the phases in terms of the quantities

$$\beta_1 = \frac{2\rho + 3 + 4\alpha}{2\rho - 1} \log 2, \quad \beta_2 = \left(1 + \frac{\alpha}{2}\right) \log 2, \quad \text{and} \quad \eta_1 = \frac{4}{\varphi}\left(1 + \frac{1}{\rho}\right),$$

as well as $\eta_2 = \eta_2(\beta)$, given by

$$\eta_2 = \frac{1}{\varphi} \min\left\{2\left(1 - \frac{\log 2}{\beta}\right), \frac{4}{1+\rho}\left(1 - \left(1 + \frac{\alpha}{2}\right)\frac{\log 2}{\beta}\right)\right\}.$$

Note that both η_1 and η_2 are proportional to $1/\varphi$, and η_2 is strictly positive when $\beta > \beta_2$.

▶ **Theorem 10** (Phase Change). *Suppose that $1/2 < \rho < \alpha - 2$ and let S_h be nondecelerating.*

(i) (At Least One Bridge) *If $\beta > \beta_1$ and $\eta > \eta_1$, then there exist positive numbers $c_1 = c_1(\beta, \eta)$ and $h_1 = h_1(\beta, \eta)$ such that $\pi(\mathsf{NB}) \leq e^{-c_1 h}$, for all $h \geq h_1$.*

(ii) (No Bridge) *If $\beta > \beta_2$ and $\eta < \eta_2(\beta)$, then there exist positive numbers $c_2 = c_2(\beta, \eta)$ and $h_2 = h_2(\beta, \eta)$ such that $\pi(\mathsf{NB}) \geq 1 - e^{-c_2 h}$, for all $h \geq h_2$.*

Since η_1 and η_2 are proportional to $1/\varphi$, an immediate corollary of Theorem 10 is that, if the scent is especially strong or especially weak, in the sense that the asymptotic growth of the column scent as $h \to \infty$ satisfies $\varphi = \omega_h(1)$ or $\varphi = o_h(1)$, then bridge formation is either likely or unlikely, independently of η.

▶ **Corollary 11.** *Suppose that $1/2 < \rho < \alpha - 2$ and S_h is nondecelerating.*

(i) (Strong Scent) *If $\beta > \beta_1$ and $\varphi = \omega_h(1)$, then there exist positive numbers $c_3 = c_3(\beta)$ and $h_3 = h_3(\beta)$ such that $\pi(\mathsf{NB}) \leq e^{-c_3 h}$, for all $h \geq h_3$.*

(ii) (Weak Scent) *If $\beta > \beta_2$ and $\varphi = o_h(1)$, then there exist positive numbers $c_4 = c_4(\beta)$ and $h_4 = h_4(\beta)$ such that $\pi(\mathsf{NB}) \geq 1 - e^{-c_4 h}$, for all $h \geq h_4$.*

Note Corollary 11(ii) holds without the assumption that S_h is nondecelerating. In fact, it holds when S_h is nonnegative because, when $\varphi = o_h(1)$, the contribution of boundary length to the Hamiltonian dominates the scent component, making fine properties of S_h irrelevant.

The proof of Theorem 10 relies on three lemmas. Informally, the first says that a subset of configurations Ω^{bad} has exponentially small probability in h if every $\sigma \in \Omega^{\text{bad}}$ can be mapped to a configuration with a Hamiltonian that is smaller by roughly $h + B(\sigma)$. The proof, which can be found in the full version of the paper [17], is similar to that of Theorem 5.

▶ **Lemma 12.** *Fix an aspect ratio α, as well as δ and ϵ, all positive numbers. Suppose that Ω^{bad} is a subset of configurations such that, for every $\sigma \in \Omega^{\text{bad}}$, there exists $\tau_\sigma \in \Omega$ satisfying*

$$H(\sigma) - H(\tau_\sigma) \geq \epsilon h + \delta B(\sigma) + O_h(1),$$

where the function implicit in $O_h(1)$ does not depend on σ. Then, for every

$$\beta > \beta_0(\delta, \epsilon) := \max\left\{\frac{\log 2}{\delta}, \frac{2\alpha \log 2}{\epsilon}\right\},$$

there exist $c = c(\beta) > 0$ and $h_0 = h_0(\beta) > 0$ such that $\pi(\Omega^{\text{bad}}) \leq e^{-ch}$ for all $h \geq h_0$.

The next lemma states that if a configuration does not form a bridge, then it is possible to identify a configuration with a Hamiltonian that is strictly smaller as a function of h and the boundary length of the original configuration.

▶ **Lemma 13.** *If $\beta > \beta_1$ and $\eta > \eta_1$, then there exist positive values $\delta_1 = \delta_1(\rho, \alpha)$ and $\epsilon_1 = \epsilon_1(\rho, \alpha)$ such that $\beta_1 \geq \beta_0(\delta_1, \epsilon_1)$ and, for every $\sigma \in \mathsf{NB}$, there exists $\tau_\sigma \in \Omega$ satisfying*

$$H(\sigma) - H(\tau_\sigma) \geq \epsilon_1 h + \delta_1 B(\sigma) + O_h(1).$$

The proof involves a multi-stage transformation of the original configuration's layer sequence (Appendix A). The same is true of the next result, which instead compares configurations that reach Λ_h to the empty configuration $\emptyset$, which has $H(\emptyset) = 0$.

▶ **Lemma 14.** *If $\beta > \beta_2$ and $\eta < \eta_2(\beta)$, then there exist positive values $\delta_2 = \delta_2(\eta, \varphi, \rho, \alpha)$ and $\epsilon_2 = \epsilon_2(\eta, \varphi, \rho, \alpha)$ such that $\beta_2 \geq \beta_0(\delta_2, \epsilon_2)$ and every $\sigma \notin \mathsf{NB}$ satisfies*

$$H(\sigma) - H(\emptyset) \geq \epsilon_2 h + \delta_2 B(\sigma) + O_h(1).$$

Theorem 10 follows directly from the preceding lemmas.

Proof of Theorem 10. Parts (i) and (ii) follow from two applications of Lemma 12, one to $\Omega^{\text{bad}} = \mathsf{NB}$ and another to $\Omega^{\text{bad}} = \mathsf{NB}^c$, which are justified by Lemmas 13 and 14. ◀

3.3 Layer sequences

The proofs of our main results rely on mappings of configurations to ones with more desirable properties (Lemmas 6, 7, 13, and 14). For example, Lemma 6 states that every configuration can be mapped to a second configuration with at least as favorable boundary length and scent components, and which has a boundary length of $O(w + h)$. The purpose of this section is to introduce layer sequences, the key technical idea underlying these lemmas, and to demonstrate their use by proving Lemma 6.

The layer sequence $\overline{N} = (n_k)_{k=1}^h$ of a (possibly non-valid) configuration $\sigma \subseteq V(\Lambda)$ counts the number $n_k = n_k(\sigma)$ of its elements in each layer of Λ:

$$n_k(\sigma) = |\sigma \cap V(\Lambda_k)|.$$

The statements that we will make in this section will apply not only to configurations in Ω, but a broader class $\overline{\Omega}$ of configurations $\sigma \subseteq V(\Lambda)$ that have layer sequences that have no fully occupied layer (k where $n_k(\sigma) = w$) after one that is not fully occupied, and no fully unoccupied layer (k where $n_k(\sigma) = 0$) before one that is not fully unoccupied:

$$\overline{\Omega} = \big\{\sigma \subseteq V(\Lambda) : (n_{k+1}(\sigma) = w \implies n_k(\sigma) = w) \text{ and } (n_k(\sigma) = 0 \implies n_{k+1}(\sigma) = 0) \text{ for all } k \in [h-1]\big\}.$$

Notably, configurations in Ω must have this property as they are simply connected, so we can say that $\Omega \subseteq \overline{\Omega}$. Referencing this property, for $0 \leq R_{bot} \leq R_{top} \leq h$, we can denote by $\mathcal{L}^{R_{top}, R_{bot}}$ the set of layer sequences that are fully occupied from layers 1 to R_{bot}, partially occupied from layers $R_{bot} + 1$ to R_{top}, and fully unoccupied for the remaining layers.

Even though the layer sequence of a configuration discards information about the configuration's boundary, we can still use it to obtain bounds on the boundary length (our definition of boundary length extends to configurations in $\overline{\Omega}$). This is the content of the next two results. To state them, we denote by σ_M the part of a configuration $\sigma \in \overline{\Omega}$ in layers 1 through M:

$$\sigma_M = \cup_{k=1}^M (\sigma \cap V(\Lambda_k)).$$

▶ **Lemma 15** (Boundary length increment). *Let $0 \le R_{bot} \le R_{top} \le h$ and let $\sigma \in \overline{\Omega}$ have layer sequence $\overline{N} \in \mathcal{L}^{R_{top},R_{bot}}$. For integers M where $\max\{R_{bot}, 1\} \le M \le \min\{R_{top} - 1, h - 2\}$, the truncation σ_{M+1} has boundary length satisfying $B(\sigma_{M+1}) \ge B(\sigma_M) + D_M$, where*

$$D_M = 2 + 2\max\{n_{M+1} - n_M + 1, 0\} + 2\max\{n_{M+1} - n_M - 1, 0\}.$$

If $M = h - 1$ instead, then we have $B(\sigma_{M+1}) = B(\sigma) \ge B(\sigma_M) + D_M - 2n_h$. Furthermore, the bound is tight in the sense that, for any $\overline{N} \in \mathcal{L}^{R_{top},R_{bot}}$, there exists $\tau \in \Omega$ (not $\overline{\Omega}$) with layer sequence $\overline{N}$ which satisfies $B(\tau_{M+1}) \le B(\tau_M) + D_M$ for every such M.

The proof of Lemma 15, which is self-contained, appears in Appendix A. Lemma 15 implies a lower bound on a configuration's boundary length in terms of its layer sequence.

▶ **Lemma 16** (Boundary length lower bounds). *Suppose that $\overline{N} \in \mathcal{L}^{R_{top},R_{bot}}$ where $0 \le R_{bot} \le R_{top} \le h$ and let $R_{min} = \max\{R_{bot}, 1\}$. For each $M \in \{R_{min}, R_{min}+1, \ldots, R_{top}\}$, we define*

$$\overline{B}_M(\overline{N}) = \overline{B}_{R_{min}}(\overline{N}) + \sum_{k=R_{min}}^{M-1} D_k, \quad \text{where} \quad \overline{B}_{R_{min}}(\overline{N}) = \begin{cases} 2n_1 + 2 & \text{if } R_{bot} = 0, \\ 2w & \text{if } R_{bot} \ge 1. \end{cases}$$

Then, for any configuration $\sigma \in \overline{\Omega}$ with layer sequence $\overline{N}$, for any $M \in \{R_{min}, R_{min} + 1, \ldots, R_{top}\}$, the truncation σ_M has boundary length satisfying

$$B(\sigma_M) \ge \begin{cases} \overline{B}_M(\overline{N}) & \text{if } M < h, \\ \overline{B}_M(\overline{N}) - 2n_h & \text{if } M = h. \end{cases}$$

Furthermore, this bound is tight in the sense that for any layer sequence $\overline{N} \in \mathcal{L}^{R_{top},R_{bot}}$, there exists a configuration $\tau \in \Omega$ with layer sequence $\overline{N}$ where $B(\tau_{R_{bot}+1}) = \overline{B}_{R_{bot}+1}(\overline{N})$ and $B(\tau_{M+1}) \le B(\tau_M) + D_M$ for each $M \in \{R_{min}, R_{min}+1, \ldots, R_{top}\}$.

Proof of Lemma 16. This follows from Lemma 15 by induction on M, with $M = R_{\min}$ as the base case. ◀

We use the two preceding facts about layer sequences to prove Lemma 6. The proof features the quantity $\overline{S}(\overline{N})$, which is the scent intensity common to all configurations with layer sequence $\overline{N}$.

Proof of Lemma 6. We aim to show that, for an arbitrary configuration $\sigma \in \Omega$, there exists $\tau \in \Omega$ such that $B(\tau) \le B(\sigma)$, $S(\tau) \ge S(\sigma)$, and $B(\tau) \le 6w + 4h$. The idea of the proof is to show that, if there are layers k_1 and k_2 such that the layer sequence $\overline{N}$ of σ satisfies

$$k_2 \ge k_1 + 2, \quad n_{k_1+1} - n_{k_1} \ge 2, \quad \text{and} \quad n_{k_2+1} - n_{k_2} \ge 2,$$

then it is possible to "promote" one of the occupied sites to a higher layer, resulting in a new layer sequence $\overline{N}^+ = (n_k^+)_{k\in[h]} \in \mathcal{L}^{R_{bot}^+,R_{top}^+}$, in such a way that $\overline{B}_h(\overline{N}^+) \le \overline{B}_h(\overline{N})$. Moving an occupied site to a higher layer clearly also ensures that $\overline{S}(\overline{N}^+) \ge \overline{S}(\overline{N})$, because S_h is nonincreasing with distance from the food. We define $\overline{N}^+$ by removing an occupied site from layer $k_1 + 1$ and adding it to layer k_2. In other words, for each $k \in [h]$, we define: $n_k^+ = n_k - 1$ if $k = k_1 + 1$; $n_k^+ = n_k + 1$ if $k = k_2$; and $n_k^+ = n_k$ otherwise. Note that, like R_{bot}, R_{bot}^+ is at most k_1.

To compute the difference between $\overline{B}_h(\overline{N}^+)$ and $\overline{B}_h(\overline{N})$, we look at the terms affected by the change. For each k, denote by D_k and D_k^+ the terms of the sums of $\overline{B}_h(\overline{N})$ and $\overline{B}_h(\overline{N}^+)$. We observe that $D_{k_1}^+ = D_{k_1} - 4$ and $D_{k_2}^+ = D_{k_2} - 4$, and the differences $D_{k_1+1}^+ - D_{k_1+1}$ and $D_{k_2-1}^+ - D_{k_2-1}$ are at most 8 when $k_1 + 1$ and $k_2 - 1$ are equal, and at most 4 when they are not. This implies that $\overline{B}_h(\overline{N}^+) \leq \overline{B}_h(\overline{N}) - 8 + 8 = \overline{B}_h(\overline{N})$.

We repeat this process of promoting occupied sites to find a layer sequence $\overline{N}^* = (n_k^*)_{k\in[h]}$ such that $\overline{B}_h(\overline{N}^*) \leq \overline{B}_h(\overline{N})$, $\overline{S}(\overline{N}^*) \geq \overline{S}(\overline{N})$, and $n_{k+1}^* - n_k^* \leq 1$ for all but at most two values of k. Moreover, if there are two such values of k, they must be consecutive integers. This implies the upper bound

$$\overline{B}_h(\overline{N}^*) \leq \overline{B}_{R_{bot}^*+1}(\overline{N}) + 4(R_{top}^* - R_{bot}^*) + 4w \leq 6w + 4h.$$

By Lemma 16, there exists $\tau \in \Omega$ with layer sequence $\overline{N}^*$ and $B(\tau) \leq \overline{B}_h(\overline{N}^*)$, hence $B(\tau) \leq 6w + 4h$. ◀

4 Conclusion

We define a simple SOPS model of bridging based on the particles' affinity for more neighbors, which results in more "robust" bridges, and a bias toward the food. We show that the emergence of a single bridge in collective systems is a statistical inevitability, requiring no central coordination. The novelty of our strategy is based on defining and analyzing a much simpler occupancy chain that abstracts the specific local dynamics of the particles and looks at the evolution of the contour indicating the system's occupied sites. We intentionally kept the occupancy chain as simple as possible to show the generality of the single bridging emergent phenomena to highlight the connections to similar statistical physics models.[2] Moreover, the occupancy chain that evolves according to a Hamiltonian defined only by contour length and scent gradient is simple enough to provide rigorous proofs using tools from statistical physics.

We see several directions for future work: First, we believe that we can extend our results to other models inspired by biological systems and beyond, such as by including ants retreat from the food after they are fed, by using a more sophisticated occupancy chain that models rafts by allowing a collection of ants to move together in a single move, or by including more specifics of man-made swarm robotics systems. Our simulations show that such variants do not significantly change our findings. Second, the occupancy chain abstracts away information about the motion of individual particles, and allows us to gain a more direct means of analyzing their collective behavior: We expect that such abstractions will help understanding the collective behavior of many other programmable matter systems. Third, occupancy Markov chains have the potential for impact in statistical physics, by allowing one to relax the need for precisely estimating surface tension of contours, potentially enabling a better formal understanding of fixed magnetization spin systems [8], and collectives arising in swarm robotics responding to directed external stimuli [12]. Lastly, while simulations suggest that both the occupancy chain and the underlying agent-based simulations converge in polynomial time, we do not have a formal bound on the mixing time of either. Future work can relax the connectivity restriction on valid configurations, which may make it easier to derive such bounds, but ant experiments suggest typically the majority do stay connected.

[2] The techniques we use also apply to other geometries, including the square lattice and other planar graphs with or without periodic boundary conditions.

It would also be helpful to derive general bounds relating the mixing times of an occupancy chain with the underlying dynamics for a general class of models, perhaps building on similar decomposition theorems [14].

References

1 M. Andrés Arroyo, S. Cannon, J.J. Daymude, D. Randall, and A.W. Richa. A stochastic approach to shortcut bridging in programmable matter. *Natural Computing*, 17(4):723–741, 2018. doi:10.1007/S11047-018-9714-X.

2 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(4):348–360, 1999.

3 S. Cannon, J.J. Daymude, C. Gökmen, D. Randall, and A.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)*, pages 54:1–54:22, 2019.

4 S. Cannon, J.J. Daymude, D. Randall, and A.W. Richa. A Markov chain algorithm for compression in self-organizing particle systems. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 279–288, 2016.

5 J. Daudelin, G. Jing, T. Tosun, M. Yim, H. Kress-Gazit, and M. Campbell. An integrated system for perception-driven autonomy with modular robots. *Science Robotics*, 3(23):eaat4983, 2018. doi:10.1126/scirobotics.aat4983.

6 J.J. Daymude, A.W. Richa, and C. Scheideler. The canonical amoebot model: Algorithms and concurrency control. *CoRR*, abs/2105.02420, 2021. arXiv:2105.02420.

7 Z. Derakhshandeh, S. Dolev, R. Gmyr, A.W. Richa, C. Scheideler, and T. Strothmann. Brief announcement: Amoebot - a new model for programmable matter. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 220–222, 2014.

8 R.L. Dobrushin, R. Kotecky, and S.B. Shlosman. *The Wulff Construction: a Global Shape from Local Interactions*. American Mathematical Society, Providence, 1992.

9 S. Greenberg, D. Randall, and A.P. Streib. Sampling biased monotonic surfaces usijng exponential metrics. *Combinatorics, Probability and Computing*, 29:672–697, 2020. doi:10.1017/S0963548320000188.

10 R. Jeanson, C. Rivault, J.L. Deneubourg, S. Blanco, R. Fournier, C. Jost, and G. Theraulaz. Self-organized aggregation in cockroaches. *Animal Behaviour*, 69(1):169–180, 2005.

11 E. Karzbrun, A.M. Tayar, V. Noireaux, and R.H. Bar-Ziv. Programmable on-chip DNA compartments as artificial cells. *Science*, 345(6198):829–832, 2014. doi:10.1126/science.1255550.

12 S. Li, B. Dutta, S. Cannon, J.J. Daymude, R. Avinery, E. Aydin, A.W. Richa, D.I. Goldman, and D. Randall. Programming active cohesive granular matter with mechanically induced phase changes. *Science Advances*, 7(17):eabe8494, 2021. doi:10.1126/sciadv.abe8494.

13 A. Lioni, C. Sauwens, G. Theraulaz, and J.-L. Deneubourg. Chain formation in Oecophylla longinoda. *Journal of Insect Behavior*, 14(5):679–696, 2001.

14 N. Madras and D. Randall. Markov chain decomposition for convergence rate analysis. *Annals of Applied Probability*, 12(2):581–606, 2002.

15 N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

16 N.J. Mlot, C.A. Tovey, and D.L. Hu. Fire ants self-assemble into waterproof rafts to survive floods. *Proceedings of the National Academy of Sciences*, 108(19):7669–7673, 2011.

17 S. Oh, J.L. Briones, J. Calvert, N. Egan, D. Randall, and A.W. Richa. Single bridge formation in self-organizing particle systems. ArXiv preprint *http://www.arxiv.org/abs/2408.10830*, 2024.

18 N.T. Ouellette and D.M. Gordon. Goals and limitations of modeling collective behavior in biological systems. *Frontiers in Physics*, 9, 2021. doi:10.3389/fphy.2021.687823.

19 C.R. Reid, M.J. Lutz, S. Powell, A.B. Kao, I.D. Couzin, and S. Garnier. Army ants dynamically adjust living bridges in response to a cost–benefit trade-off. *Proceedings of the National Academy of Sciences*, 112(49):15113–15118, 2015.

20 E. Şahin. Swarm robotics: From sources of inspiration to domains of application. In *Swarm Robotics*, pages 10–20, 2005.

21 W. Savoie, S. Cannon, J.J. Daymude, R. Warkentin, S. Li, A.W. Richa, D. Randall, and D. I. Goldman. Phototactic supersmarticles. *Artificial Life and Robotics*, 23(4):459–468, 2018. doi:10.1007/S10015-018-0473-7.

22 B. Wei, M. Dai, and P. Yin. Complex shapes self-assembled from single-stranded DNA tiles. *Nature*, 485(7400):623–626, 2012. doi:10.1038/nature11075.

23 A. Yadav. Stochastic maze solving under the geometric amoebot model. Master's thesis, Rutgers University, 2021. URL: https://rucore.libraries.rutgers.edu/rutgers-lib/65707/.

24 H. Zeng, J. Briones, R. Avinery, S. Li, A. Richa, D. Goldman, and T. Sasaki. Fire ant pontoon bridge: a self-assembled dynamic functional structure. In *Integrative and Comparative Biology*, volume 62, pages S341–S342, 2023.

A Additional details of the technical lemmas

In this appendix, we further develop the properties and applications of layer sequences, introduced in Section 3.3. First, in Section A.1, we collect the proofs of some inputs to the proof of Theorem 5. Then, in Section A.2, we discuss the main ideas behind the inputs to the proof of Theorem 10.

A.1 Further inputs to Theorem 5

Proof of Lemma 15. For $\sigma \in \overline{\Omega}$, we look at the change in total boundary length when we add layer $M+1$ of σ to σ_M. Let C_{M+1} be the number of components in the restriction of σ to only layer $M+1$.

As layer $M+1$ is not completely filled, there are exactly $n_{M+1} - C_{M+1}$ edges between occupied sites within layer $M+1$. Thus, the number of edges of Λ with exactly one endpoint at an occupied site in layer $M+1$ is equal to $6n_{M+1} - 2(n_{M+1} - C_{M+1}) = 4n_{M+1} + 2C_{M+1}$. Some of the edges are shared with occupied sites on layer M. Denoting by E_{M+1} the number of edges between layers M and $M+1$, we can see that adding layer $M+1$ to σ_M adds $4n_{M+1} + 2C_{M+1} - 2E_{M+1}$ to the boundary length of σ_{M+1}.

We compute an upper bound for E_{M+1} by counting the number of occupied sites in layer $M+1$ with at least one and with at least two edges to layer M respectively. Layer M, with C_M components, provides $n_M + C_M$ occupied sites in layer $M+1$ with at least one edge to layer M, and $n_M - C_M$ occupied sites in layer $M+1$ with two edges to layer M. A site in layer $M+1$ cannot have more than two edges to layer M. This gives an upper bound of $\min\{n_M + c_M, n_{M+1}\}$ for the former, and $\min\{n_M - c_M, n_{M+1}\}$ for the latter. This gives us the following lower bound:

$$\begin{aligned}
&B(\sigma_{M+1}) - B(\sigma_M) \\
&\quad \geq 4n_{M+1} + 2C_{M+1} - 2\min\{n_M + C_M, n_{M+1}\} - 2\min\{n_M - C_M, n_{M+1}\} \\
&\quad \geq 2C_{M+1} + 2\max\{n_{M+1} - n_M - C_M, 0\} + 2\max\{n_{M+1} - n_M + C_M, 0\} \\
&\quad \geq 2 + 2\max\{n_{M+1} - n_M - 1, 0\} + 2\max\{n_{M+1} - n_M + 1, 0\}.
\end{aligned}$$

The last inequality holds because $C_{M+1} \geq 1$ and $\max\{n_{M+1} - n_M - C_M, 0\} + \max\{n_{M+1} - n_M + C_M, 0\} \geq \max\{n_{M+1} - n_M - 1, 0\} + \max\{n_{M+1} - n_M + 1, 0\}$ for any value of $n_{M+1} - n_M$,

as long as $C_M \geq 1$ (which is true as $\sigma \in \overline{\Omega}$). To show the second statement of the lemma, we can construct σ layer by layer, where each layer from $\max\{1, R_{bot}\}$ to R_{top} has exactly one component and M_{M+1} is maximized for each value of M. ◀

Proof of Lemma 7. In this proof, we will use $(x)_+$ to denote $\max\{x, 0\}$ for $x \in \mathbb{R}$. Fix $\epsilon \geq 1/h$ and a configuration $\sigma \in \mathsf{MB}_\epsilon$. By definition (4), there is an $a \in (0, 1)$ such that σ has multiple $(a, a+\epsilon)$-bridges. Set $M = \lfloor ah \rfloor$, and let $\Lambda_{\geq M}$ represent the set of sites on layers M to h. We can divide the elements of σ within $\Lambda_{\geq M}$ into connected components $V_1, V_2, \ldots, V_P$ for some integer P. Denote by V_{base} the elements of σ not in $\Lambda_{\geq M}$. For each $j \in \{1, 2, \ldots, P\}$, we denote $\sigma^{(j)} = V_j \cup V_{\text{base}}$ which we note will be a configuration in $\overline{\Omega}$, and let $\overline{N}^{(j)} = (n_k^{(j)})_{k \in \{0,1,\ldots,h\}}$ be its layer sequence.

By applying Lemma 15, the increase in boundary length when adding the sites V_j to the restriction of σ to the layers $1, 2, \ldots, M-1$ must be at least $\overline{B}^{(j)}$, where

$$\overline{B}^{(j)} = \sum_{k=M}^{R_{top}^{(j)}-1} \left(2 + 2\left(n_{k+1}^{(j)} - n_k^{(j)} + 1\right)_+ + 2\left(n_{k+1}^{(j)} - n_k^{(j)} - 1\right)_+\right),$$

and where $R_{top}^{(j)}$ denotes the topmost non-empty layer of $\sigma^{(j)}$. We can thus give the following lower bound for the boundary length of σ:

$$B(\sigma) \geq \overline{B}_M(\overline{N}) + \sum_{j=1}^{P} \overline{B}^{(j)}.$$

Denoting by $\overline{N} = (n_k)_{k \in \{0,1,\ldots,h\}}$ the layer sequence of σ itself, by Lemma 16, there exists a configuration τ with layer sequence $\overline{N}$ where:

$$B(\tau) = \overline{B}_M(\overline{N}) + \sum_{k=M}^{R_{top}-1} \left(2 + 2\left(n_{k+1} - n_k + 1\right)_+ + 2\left(n_{k+1} - n_k - 1\right)_+\right).$$

Without loss of generality, we may assume that $R_{top} = R_{top}^{(1)} \geq R_{top}^{(2)} \geq \ldots \geq R_{top}^{(P)}$.

$$\begin{aligned}
B(\sigma) - B(\tau) &\geq \sum_{j=1}^{P} \overline{B}^{(j)} - \sum_{k=M}^{R_{top}} \left(2 + 2\left(n_{k+1} - n_k + 1\right)_+ + 2\left(n_{k+1} - n_k - 1\right)_+\right) \\
&= -\sum_{j=2}^{P} \sum_{k=M}^{R_{top}^{(j)}} 2 + 2 \sum_{k=M}^{R_{top}^{(j)}} \left(\sum_{\substack{j \in [P] \\ k \leq R_{top}^{(j)}-1}} \left(n_{k+1}^{(j)} - n_k^{(j)} + 1\right)_+ - \left(n_{k+1} - n_k + 1\right)_+ \right) \\
&\qquad + 2 \sum_{k=M}^{R_{top}^{(j)}} \left(\sum_{\substack{j \in [P] \\ k \leq R_{top}^{(j)}-1}} \left(n_{k+1}^{(j)} - n_k^{(j)} - 1\right)_+ - \left(n_{k+1} - n_k - 1\right)_+ \right).
\end{aligned}$$

The second and third terms of the above expression are non-negative as for $I \in \{1, -1\}$, we

have

$$\sum_{\substack{j\in[P]\\ k\le R_{top}^{(j)}-1}} \left(n_{k+1}^{(j)} - n_k^{(j)} + I\right)_+ - (n_{k+1} - n_k + I)_+$$
$$= \sum_{\substack{j\in[P]\\ k\le R_{top}^{(j)}}} \left(n_{k+1}^{(j)} - n_k^{(j)} + I\right)_+ - \Big(\sum_{\substack{j\in[P]\\ k\le R_{top}^{(j)}}} (n_{k+1} - n_k) + I\Big)_+ \ge 0.$$

This allows us to conclude that

$$B(\sigma) - B(\tau) \ge \sum_{j=2}^{P}\sum_{k=M}^{R_{top}^{(j)}} 2 = 2\sum_{j=2}^{P}(R_{top}^{(j)} - M + 1). \tag{10}$$

If σ has multiple $(a, a+\epsilon)$-bridges, we must have $P \ge 2$ and $R_{top}^{(2)} - M + 1 \ge \epsilon h$, and so according to (10) we must have $B(\sigma) \ge B(\tau) + 2\epsilon h$. ◀

A.2 Further inputs to Theorem 10

In this section, we present the main ideas behind Lemmas 13 and 14, which we prove in the full version of the paper [17]. These lemmas arise from the analysis of a transformation of layer sequences which gives a significant improvement (more negative) in the Hamiltonian when a configuration without a bridge to the food is extended to reach the food, so long as the scent parameter η is sufficiently large.

Let σ be an arbitrary configuration with layer sequence $\overline{N} = (n_k)_{k\in[h]}$. Suppose that $\overline{N} \in \mathcal{L}^{0,D}$, where $D \le h-1$ is the highest occupied layer in σ. We will transform $\overline{N}$ into a new layer sequence $\overline{N}^{\text{post}}$ in $\mathcal{L}^{0,h}$ (i.e. it reaches Λ_h). We then show that with a sufficiently large value of η (or φ), the increase in scent intensity going from $\overline{N}$ to $\overline{N}^{\text{post}}$ will more than compensate for the potential increase in boundary length. However, the transformation from $\overline{N}$ to $\overline{N}^{\text{post}}$ is complex, and we will thus split it into multiple steps – pre-process, the main transformation, and post-process. The layer sequences after each of these steps are denoted $\overline{N}^{\text{pre}}$, $\overline{N}^{\text{mid}}$ and $\overline{N}^{\text{post}}$ respectively.

To understand the following transformations from $\overline{N}$ to $\overline{N}^{\text{pre}}$ to $\overline{N}^{\text{mid}}$ to $\overline{N}^{\text{post}}$, we can visualize layer sequences as "right-justified" configurations in Λ (see Figure 3). The particles of these right-justified configurations can be grouped into columns, where the site on the k^{th} layer of column j (counted from the right) is filled if and only if its layer sequence has at least j particles in layer k.

The pre-processing step, which creates an intermediate layer sequence $\overline{N}^{\text{pre}} = (n_k^{\text{pre}})_{k\in[h]} \in \mathcal{L}^{0,h}$, is solely to account for specific edge cases in the eventual transformation to $\overline{N}^{\text{post}}$. To define $\overline{N}^{\text{pre}}$ starting from $\overline{N}$, we first add in the unused $n - |\sigma|$ particles to the layer sequence. We denote $u = w - \max_{k\in[h]}\{n_k\} - 1$, the number of empty columns that we can fill with unused particles. We thus take $\min\left\{u, \lfloor\frac{n-|\sigma|}{h}\rfloor\right\}$ empty columns and fill them with unused particles (which is equivalent to increasing each entry of the layer sequence by said amount of columns). The next step is to guarantee at least one completely filled column. If there are no completely filled columns even after adding the unused particles, $h - D$ particles are taken one at the time from the highest layer with at least two particles, and placed to form a single column of particles. This transformation is illustrated in Figure 3c.

The particles of the layer sequence $\overline{N}^{\text{pre}}$ can be divided into three layer sequences $(n_k^{\text{full}})_{k\in[h]}$, $(n_k^{\perp})_{k\in[h]}$ and $(n_k^{\top})_{k\in[h]}$ by column (so that $n_k^{\text{pre}} = n_k^{\text{full}} + n_k^{\perp} + n_k^{\top}$ for all $k \in [h]$).

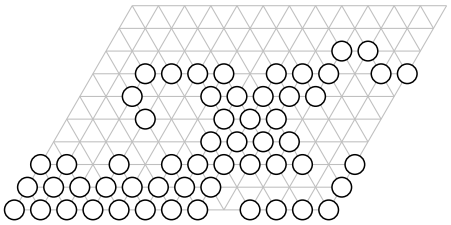

(a) Original configuration σ with some number of unused particles (not shown).

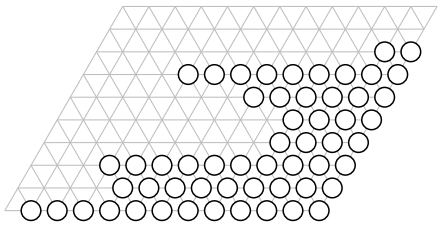

(b) Representation of layer sequence $\overline{N}$ (right-justified σ) with unused particles not shown.

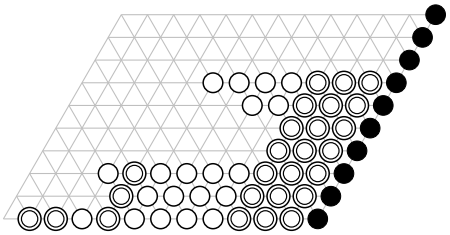

(c) Representation of layer sequence $\overline{N}^{\text{pre}}$. There were no empty columns in $\overline{N}$, so all unused particles remain unused, while 2 particles from the topmost layers are moved to form a single column to the top. Particles from $(n_k^{\text{full}})_{k\in[h]}$, $(n_k^{\perp})_{k\in[h]}$ and $(n_k^{\top})_{k\in[h]}$ represented by black circles, double circles and single circles respectively.

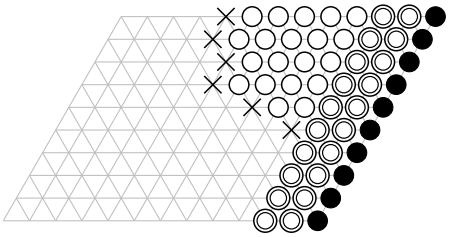

(d) Representation of layer sequence $\overline{N}^{\text{mid}}$ (excluding the crosses). Particles from $(n_k^{\top})_{k\in[h]}$ are shifted upward, and two full columns are created from the 26 particles in $(n_k^{\perp})_{k\in[h]}$, with 6 particles from $(n_k^{\perp})_{k\in[h]}$ unused. These 6 particles are added back to the positions marked with crosses to form the layer sequence $\overline{N}^{\text{post}}$.

Figure 3 Transformation from the layer sequence $\overline{N}$ to $\overline{N}^{\text{post}}$, with the layer sequences represented as right-justified configurations.

The first layer sequence $(n_k^{\text{full}})_{k\in[h]}$ is represents the particles belonging to the $J^{\text{full}} = n_h^{\text{pre}} \geq 1$ full columns of the right-justified configuration. The remaining columns are then identified as either *bottom-supported* or *non-bottom-supported*. The bottom-supported columns are those which have some value $k' \in \{0, 1, 2, \ldots, h-1\}$ where layers $\{1, 2, \ldots, k'\}$ in the column are occupied and the remaining layers are unoccupied (an alternative definition is having no unoccupied site below an occupied site). Particles from bottom-supported columns are put into $(n_k^{\perp})_{k\in[h]}$, while those from non-bottom-support columns go into $(n_k^{\top})_{k\in[h]}$. Alternatively, for each value of k in $[h]$ we can define formally

- $n_k^{\text{full}} = J^{\text{full}}$,
- $n_k^{\perp} = \left|\left\{j \in [w] : j > J^{\text{full}} \text{ and } n_{k'+1}^{\text{pre}} \geq j \implies n_{k'}^{\text{pre}} \geq j \text{ for all } k' \in [h-1]\right\}\right|$, and
- $n_k^{\top} = n_k^{\text{pre}} - n_k^{\perp} - n_k^{\text{full}}$.

To construct $\overline{N}^{\text{mid}}$, we apply different transformations to the particles corresponding to each of these layer sequences to construct a new configuration σ^{mid}. The particles in $(n_k^{\perp})_{k\in[h]}$ are transformed into $J^{\perp} = \lfloor \frac{1}{h} \sum_{k=1}^{h} n_k^{\perp} \rfloor$ full columns while the particles in $(n_k^{\top})_{k\in[h]}$ have their particles shifted to the tops of their respective columns. These non-bottom-supported columns are placed directly next to the $J^{\text{full}} + J^{\perp}$ filled columns in the right-justified configuration σ^{mid} (see Figure 3d). Note that there may be up to $h-1$ unused particles from $(n_k^{\perp})_{k\in[h]}$ after this transformation. The layer sequence of this new configuration will be denoted by $\overline{N}^{\text{mid}} = (n_k^{\text{mid}})_{k\in[h]}$.

To describe this transformation more precisely in terms of layer sequences, for each $i \in [w]$, we denote $\mathcal{C}_i^{\top} = |\{k \in [h] : n_k^{\top} \geq i\}|$, the number of particles in the i^{th} non-bottom-

supported column. With J^{full} and $J^{\perp}$ defined as before, we can then define the layer sequence $\overline{N}^{\text{mid}} = (n_k^{\text{mid}})_{k\in[h]}$, where for each $k \in [h]$,

$$n_k^{\text{mid}} = J^{\text{full}} + J^{\perp} + \left|\left\{i \in [w] : \mathcal{C}_i^{\top} \geq h - k + 1\right\}\right|.$$

We note that this is a valid transformation (into a layer sequence $\overline{N}^{\text{mid}} \in \mathcal{L}^{0,h}$) as for each $k \in [h]$, we know that $J^{\perp} \leq \max_{k'}\{n_{k'}^{\perp}\}$ and $\left|\left\{i \in [w] : \mathcal{C}_i^{\top} \geq h - k + 1\right\}\right| \leq \max_{k'}\{n_{k'}^{\top}\}$, so we must have $n_k^{\text{mid}} \leq n_k^{\text{pre}} \leq w - 1$.

We include one final step in the transformation to ensure that the change in scent intensity from $\overline{N}^{\text{pre}}$ to $\overline{N}^{\text{post}}$ is non-negative. In particular, in the transformation from $\overline{N}^{\text{pre}}$ to $\overline{N}^{\text{post}}$, for the sake of simplicity, we had transformed the particles from $(n_k^{\perp})_{k\in[h]}$ into $J^{\perp} = \lfloor \frac{1}{h}\sum_{k=1}^{h} n_k^{\perp} \rfloor$ full columns, leaving $m \leq h - 1$ particles unused. As a final post-processing step, we can add these m particles back to the final layer sequence, by adding one particle each to the topmost m rows that have less than $w - 1$ particles (note that the final layer sequence is non-decreasing). This gives a positive net change in scent intensity for the particles from $(n_k^{\perp})_{k\in[h]}$, while potentially increasing the value of $\overline{B}_h$ by at most $O_h(1)$ (using the definition of $\overline{B}_h$ in Lemma 16).

Change in scent intensity. We compute the overall change in the two transformations to show Lemma 17. It is important to note that in our analysis, we do not assume that φ is constant with respect to h. Note that for this and the following lemmas, we will only give the brief ideas for the proofs; the remaining details are presented in [17]. Recall that $\overline{S}(\overline{N}')$ denotes the scent intensity common to all configurations with layer sequence $\overline{N}'$.

▶ **Lemma 17** (Scent Intensity). *The changes in scent intensities in the transformations from $\overline{N}$ to $\overline{N}^{pre}$ and from $\overline{N}^{pre}$ to $\overline{N}^{post}$ are non-negative and have the following lower bounds:*

$$\overline{S}(\overline{N}^{pre}) - \overline{S}(\overline{N}) \geq \varphi \cdot \min\left\{u, \left\lfloor \frac{n - |\sigma|}{h} \right\rfloor\right\}$$

$$\overline{S}(\overline{N}^{post}) - \overline{S}(\overline{N}^{pre}) \geq |\sigma|\frac{\varphi}{h}\left(1 - \frac{D}{h}\right) + \varphi \cdot O_h(1),$$

where $u = w - \max_{k\in[h]}\{n_k\} - 1$ and $O_h(1)$ represents a function that does not depend on the specific choice of starting configuration σ.

For the transformation from $\overline{N}$ to $\overline{N}^{\text{pre}}$ in the proof of Lemma 17, the number of full columns added to the $\overline{N}$ is $\min\left\{u, \left\lfloor \frac{n-|\sigma|}{h} \right\rfloor\right\}$, and neither the adding of full columns nor the shifting of $h - D$ particles to the top $h - D$ rows can decrease the total scent intensity. On the other hand, the transformation from $\overline{N}^{\text{pre}}$ to $\overline{N}^{\text{post}}$ requires a more involved analysis, to show that for particles from both bottom-supported columns and non-bottom-supported columns, the average increase in the total scent intensity per particle has a lower bound of $\frac{\varphi}{h}\left(1 - \frac{D}{h}\right)$, with a small error term.

Change in boundary length. Instead of computing the boundary lengths for configurations directly, we compare the (tight) boundary length lower bounds that are defined in Lemma 16 between the initial and final layer sequences $\overline{N}$ and $\overline{N}^{\text{post}}$.

▶ **Lemma 18** (Boundary Length). *In the transformation from $\overline{N} \in \mathcal{L}^{0,D}$ to $\overline{N}^{post} \in \mathcal{L}^{0,h}$, we have:*

$$\overline{B}_h(\overline{N}^{post}) - \overline{B}_D(\overline{N}) \leq 4(h - D) - 2\max\left\{w - u - D, \left(\frac{|\sigma|}{h} - \frac{h}{2}\right)\right\} + O_h(1),$$

where $O_h(1)$ represents a function that does not depend on the specific choice of starting configuration σ. Furthermore, by Lemma 16 there exists a configuration τ with layer sequence $\overline{N}^{post}$ such that $B(\tau) - B(\sigma) \leq \overline{B}_h(\overline{N}^{post}) - \overline{B}_D(\overline{N})$, and thus has the same upper bound.

The proof of Lemma 18 focuses almost entirely on the comparison between layer sequences $\overline{N}$ and $\overline{N}^{\text{mid}}$, by noting that the post-processing step produces an insignificant change in boundary length.

Change in the Hamiltonian. Finally, putting together the changes in scent intensity (Lemma 17) and boundary length (Lemma 18) going form $\overline{N}$ to $\overline{N}^{\text{post}}$, we can compute a lower bound for the change in Hamiltonian. There are two cases in the proof of Lemma 13, roughly corresponding to whether or not there were enough empty columns u in $\overline{N}$ to fit the majority of the unused particles in the transformation from $\overline{N}$ to $\overline{N}^{\text{pre}}$. In the case where u is small ($u < \lfloor \frac{n-|\sigma|}{h} \rfloor$), the boundary length of the initial configuration or layer sequence will be close to w, and is the primary contributor to the change in Hamiltonian. In the case where u is large ($u \geq \lfloor \frac{n-|\sigma|}{h} \rfloor$), the increase in scent intensity without additional boundary length from being able to construct full columns of particles makes it undesirable for particles to remain unused in the configuration.