

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

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

Joshua J. Daymude†
Arizona State University
Tempe, AZ, USA
jdaymude@asu.edu

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

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

Andréa W. Richa¶
Arizona State University
Tempe, AZ, USA
aricha@asu.edu

ABSTRACT

We investigate stochastic, distributed algorithms that can accomplish separation and integration behaviors in *self-organizing particle systems*, an abstraction of programmable matter. These particle systems are composed of individual computational units known as *particles* that have limited memory, strictly local communication abilities, and modest computational power, and which collectively solve system-wide problems of movement and coordination. In this work, we extend the usual notion of a particle system to treat *heterogeneous* systems by considering particles of different colors. We present a fully distributed, asynchronous, stochastic algorithm for *separation*, where the particle system self-organizes into segregated color classes using only local information about each particle’s preference for being near others of the same color. Conversely, by simply changing the particles’ preferences, the color classes become well-integrated. We rigorously analyze the convergence of our distributed, stochastic algorithm and prove that under certain conditions separation occurs. We also present simulations demonstrating our algorithm achieves both separation and integration.

CCS CONCEPTS

• **Theory of computation** → **Distributed algorithms**; *Random walks and Markov chains*; • **Computing methodologies** → *Self-organization*;

*This material is based on work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. NSF DGE-1650044. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

†Supported in part by NSF Awards CCF-1637393 and CCF-1733680.

‡Supported in part by NSF Award CCF-1733812.

§Supported in part by NSF Awards CCF-1637031 and CCF-1733812.

¶Supported in part by NSF Awards CCF-1637393 and CCF-1733680.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PODC '18, July 23–27, 2018, Egham, United Kingdom

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5795-1/18/07.

<https://doi.org/10.1145/3212734.3212792>

KEYWORDS

Programmable matter; self-organization; separation; distributed algorithms; stochastic algorithms; Markov chains

ACM Reference Format:

Sarah Cannon, Joshua J. Daymude, Cem Gokmen, Dana Randall, and Andréa W. Richa. 2018. Brief Announcement: A Local Stochastic Algorithm for Separation in Heterogeneous Self-Organizing Particle Systems. In *PODC '18: ACM Symposium on Principles of Distributed Computing, July 23–27, 2018, Egham, United Kingdom*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3212734.3212792>

1 INTRODUCTION

Examples of heterogeneous entities separating and integrating exist at many scales, from molecules exhibiting attraction and repulsion to inherent human biases that influence how we form and maintain social groups. Another example is species such as ants co-mingling peacefully when resources are plentiful but prioritizing the survival of their own colony when resources are scarce. This fundamental behavior of separation or integration in response to environmental stimuli spans remarkably diverse disciplines.

We focus on *programmable matter*, a physical material or substance that can intelligently respond to user input or environmental stimuli by changing its physical properties to achieve a goal. We abstractly envision programmable matter as a *self-organizing particle system*, an ensemble of simple active computational *particles* that individually execute distributed, local, asynchronous algorithms to cooperatively achieve macro-scale objectives. Here we consider *heterogeneous* particle systems — where particles have immutable *colors* — and seek local, distributed algorithms which result in *separation* of color classes.

To develop distributed algorithms for separation, we use concepts from stochastic processes. Of particular relevance is the Schelling model [11, 12] which explores how micro-motives can induce macro-phenomena such as racial segregation in residential neighborhoods. Stochastic dynamics for the Ising model of ferromagnetism from statistical physics [13] exhibit a similar dependence of global behavior on a single parameter controlling local preferences. Our work harnesses this interplay between local preferences and global behavior to develop a stochastic, distributed, asynchronous algorithm that provably accomplishes separation.




Figure 1: (a) A section of the triangular lattice. (b) Contracted and expanded particles. The former are shown as black circles, and the latter as two black circles joined by a black line.

1.1 The Amoebot Model

In the *amoebot model*, programmable matter consists of individual, homogeneous computational elements called *particles*. A full description of the model is available online [6]. In the *geometric amoebot model*, we further assume the underlying geometry is the infinite triangular lattice (Fig. 1a). Each particle occupies a single node (it is *contracted*) or a pair of adjacent nodes (it is *expanded*); see Fig. 1b. Particles move via a series of *expansions* and *contractions*: a contracted particle can expand into an unoccupied adjacent node, and completes its movement by contracting to once again occupy a single node. Two particles occupying adjacent nodes are *neighbors*. A particle communicates only with its neighbors, has constant-size memory, and does not have any global information such as a coordinate system, orientation, or compass. A particle system is *connected* if the subgraph of the triangular lattice induced by its occupied nodes is connected. A *hole* of a particle system is a maximal finite, connected component of unoccupied nodes.

1.2 Systems of Heterogeneous Particles

We consider a particle system composed of n heterogeneous particles, generalizing previous work where particles were identical and indistinguishable [1, 5]¹. We model heterogeneity by assuming each particle P has a color $c(P) \in \{c_1, \dots, c_k\}$ in its memory, visible to itself and its neighbors; we will assume $k = 2$. These colors can represent anything from differences in equipment between robots in multi-robot systems to demographic diversity in human communities. If particles P and Q are neighbors we say they are joined by an *edge*; this edge is *homogeneous* if $c(P) = c(Q)$ and *heterogeneous* otherwise.

We define a *swap move* under the amoebot model that enables adjacent particles of different colors to switch places. For two neighboring contracted particles P and Q , either P or Q can initiate a swap in one atomic action, which can be implemented as follows: P reads $x \leftarrow c(Q)$ from the memory of Q , overwrites $c(Q) \leftarrow c(P)$ in the memory of Q , and finally updates $c(P) \leftarrow x$. Implementing a swap as an exchange of in-memory attributes is purely for modeling convenience. When individuals have immutable “color” (e.g., ants from different colonies or robots with different hardware), swaps could be realized by a coordinated movement. Adding this natural swap move enables faster convergence of our algorithms in practice, but is not necessary for our results.

¹A particle system with particles of different functional capabilities was considered in [7]; our particles all have the same capabilities regardless of color.

1.3 The Stochastic Approach

The stochastic (Markov chain) approach to self-organizing particle systems was introduced in [5] and further validated in [1]. Background on Markov chains can be found in standard textbooks (e.g., [9]) or in previous work [1, 5]. This approach is motivated by work in statistical physics that investigates the local micro-behavior causes of global macroscopic phenomena (e.g., [2, 3, 10]). Like a spring relaxing, physical systems favor configurations σ that minimize energy, determined by a *Hamiltonian* $H(\sigma)$. Each configuration is given a weight from the *Gibbs distribution*: $w(\sigma) = e^{-B \cdot H(\sigma)} / Z$, where B is inverse temperature and $Z = \sum_{\sigma} e^{-B \cdot H(\sigma)}$ is the normalizing constant (or *partition function*).

To achieve separation, we define a Hamiltonian where particle configurations with many edges and large monochromatic clusters have the lowest values. We then get Gibbs distribution that simplifies to $w(\sigma) = \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)} / Z$, where $e(\sigma)$ is the number of edges and $a(\sigma)$ is the number of homogeneous edges in σ . Larger λ favors more compressed configurations, while for smaller λ the opposite is true, just as in [5]. We expect γ to induce separation when large and integration when small.

Our main contribution is a Markov chain \mathcal{M} whose stationary distribution π is exactly this Gibbs distribution. We run \mathcal{M} indefinitely; once we reach π , we continue moving among different configurations but remain at this desirable distribution. Showing poorly separated configurations are exponentially unlikely in π is our main technical result. We carefully design \mathcal{M} using a *Metropolis filter* [8], ensuring the probabilities of all transitions of \mathcal{M} can be calculated with local information, which is necessary for translating \mathcal{M} into a distributed algorithm.

2 ALGORITHM FOR SEPARATION

Starting at any connected, hole-free particle configuration, Markov chain \mathcal{M} (Algorithm 1) for separation ensures the particle system remains connected and hole-free throughout its execution. We use the following notation. For a location ℓ , let $N_i(\ell)$ denote the set of particles of color c_i adjacent to location ℓ . For neighboring locations ℓ and ℓ' , let $N_i(\ell \cup \ell')$ be the set $N_i(\ell) \cup N_i(\ell')$, excluding particles occupying ℓ and ℓ' . When ignoring color, let $N(\ell) = \bigcup_i N_i(\ell)$; define $N(\ell \cup \ell')$ analogously. Let $\mathbb{S} = N(\ell) \cap N(\ell')$ be the set of particles adjacent to both locations. The following locally-checkable properties ensure the particles stay connected and hole-free.

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

PROPERTY 2. $|\mathbb{S}| = 0$, and both $N(\ell) \setminus \{\ell'\}$ and $N(\ell') \setminus \{\ell\}$ are nonempty and connected.

To translate \mathcal{M} into a fully local, distributed, asynchronous algorithm \mathcal{A} run by each particle concurrently, we decompose the steps of \mathcal{M} into individual particle activations, in which a single particle performs some computation and at most one movement [6]. This decouples a particle’s expansion and contraction (Steps 4–7) and P and Q ’s coordinated swap (Steps 9–11) each into two particle activations. We use flag-locking mechanisms similar to [5] to ensure consistent snapshots of particle neighborhoods; details are omitted.

We prove the following results about \mathcal{M} , which then also hold for \mathcal{A} . Proofs can be found in [4].




Figure 2: From left to right, a 2-color particle system after 0; 50,000; 1,050,000; 17,050,000; and 68,250,000 iterations of \mathcal{M} with $\lambda = 4$ and $\gamma = 4$.

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

- 1: Choose a particle P uniformly at random from all n particles; let c_i be its color and ℓ its location.
- 2: Choose a neighboring location ℓ' and $q \in (0, 1)$ uniformly at random.
- 3: **if** ℓ' is unoccupied **then**
- 4: P expands to occupy both ℓ and ℓ' .
- 5: **if** (i) ℓ and ℓ' satisfy Property 1 or 2 and (ii) $q < \lambda^{|N(\ell')| - |N(\ell)|} \cdot \gamma^{|N_i(\ell')| - |N_i(\ell)|}$ **then**
- 6: P contracts to ℓ' .
- 7: **else** P contracts back to ℓ .
- 8: **else if** ℓ' is occupied by particle Q of color c_j **then**
- 9: P calculates $|N_i(\ell)|$ and $|N_j(\ell) \setminus \{Q\}|$ and sends them to Q .
- 10: Q calculates $|N_i(\ell') \setminus \{P\}|$ and $|N_j(\ell')|$.
- 11: **if** $q < \gamma^{|N_i(\ell') \setminus \{P\}| - |N_i(\ell)| + |N_j(\ell) \setminus \{Q\}| - |N_j(\ell')|}$ **then** Q swaps with P .

LEMMA 2.1. *The particle system remains connected and hole-free throughout the execution of \mathcal{M} .*

LEMMA 2.2. *\mathcal{M} is ergodic, converging to unique stationary distribution π given by $\pi(\sigma) = \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)} / Z$, where $Z = \sum_{\sigma} \lambda^{e(\sigma)} \cdot \gamma^{a(\sigma)}$.*

Informally, a configuration with two colors is *separated* if there is a set R of particles such that most particles in R have color c_1 , most particles in its complement \bar{R} have color c_2 , and the boundary between R and \bar{R} is small. If this holds, R and \bar{R} are *clusters*. Formally, a configuration is (β, δ) -clustered, for $\beta > 0$ and $\delta < 1/2$, if there are at most $\delta|R|$ particles of color c_2 in R , at most $\delta|\bar{R}|$ particles of color c_1 in \bar{R} , and the boundary between R and \bar{R} is of size at most $\beta\sqrt{n}$, where n is the number of particles.

THEOREM 2.3 (INFORMAL). *Among the particle system configurations within an α factor of the minimum possible perimeter, for any $\beta > 4\alpha$ and $\delta < 1/2$, the probability \mathcal{M} is not in a (β, δ) -clustered configuration at stationarity is exponentially small when λ , γ , and the number of particles are large enough.*

3 SIMULATIONS

We simulated \mathcal{M} on a system of 100 particles with two colors. Fig. 2 shows the progression of \mathcal{M} over time with parameters $\lambda = 4$ and $\gamma = 4$, the regime where we expect compression and separation. Fig. 3 shows the results of running \mathcal{M} from the same initial configuration for the same number of iterations, varying only λ and γ ; we observe four distinct phases: expanded-integrated, expanded-separated, compressed-integrated, and compressed-separated.

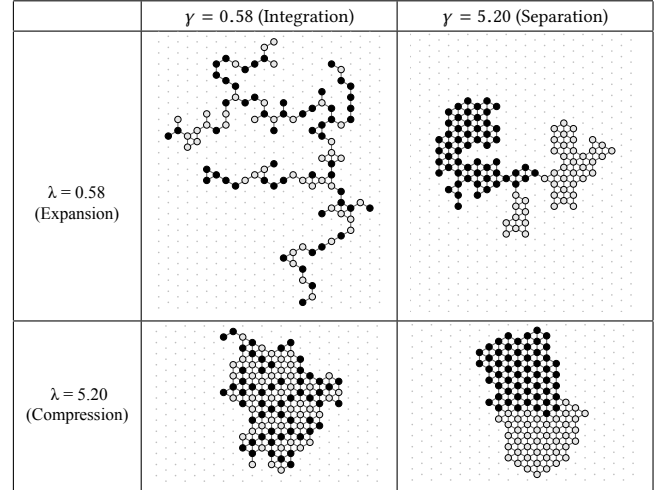


Figure 3: A 2-color particle system starting in the leftmost configuration of Fig. 2 after 50 million iterations of \mathcal{M} for various λ and γ .

REFERENCES

- [1] M. Andrés Arroyo, S. Cannon, J.J. Daymude, D. Randall, and A.W. Richa. 2017. A Stochastic Approach to Shortcut Bridging in Programmable Matter. In *DNA Computing and Molecular Programming – 23rd International Conference (DNA23)*. 122–138. A full and updated version is at <https://arxiv.org/abs/1709.02425>.
- [2] R. J. Baxter, I. G. Enting, and S. K. Tsang. 1980. Hard-square lattice gas. *Journal of Statistical Physics* 22, 4 (01 Apr 1980), 465–489.
- [3] A. Blanca, D. Galvin, D. Randall, and P. Tetali. 2013. Phase Coexistence and Slow Mixing for the Hard-Core Model on \mathbb{Z}^2 . In *17th International Workshop on Randomization and Computation (RANDOM '13)*. 379–394.
- [4] S. Cannon, J.J. Daymude, C. Gokmen, D. Randall, and A.W. Richa. 2018. A Local Stochastic Algorithm for Separation in Heterogeneous Self-Organizing Particle Systems. (2018). Available at <https://arxiv.org/abs/1805.04599>.
- [5] S. Cannon, J.J. Daymude, D. Randall, and A.W. Richa. 2016. A Markov Chain Algorithm for Compression in Self-Organizing Particle Systems. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC '16)*. 279–288.
- [6] J.J. Daymude, A.W. Richa, and C. Scheidele. 2017. The Amoebot Model. (2017). Available online at <https://sops.engineering.asu.edu/sops/amoebot>.
- [7] G.A. Di Luna, P. Flocchini, G. Prencipe, N. Santoro, and G. Viglietta. 2018. Line Recovery by Programmable Particles. In *Proceedings of the 19th International Conference on Distributed Computing and Networking (ICDCN '18)*. 4:1–4:10.
- [8] W. K. Hastings. 1970. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* 57 (1970), 97–109.
- [9] D. Levin, Y. Peres, and E. Wilmer. 2009. *Markov chains and mixing times*. Providence, R.I. American Mathematical Society.
- [10] R. Restrepo, J. Shin, P. Tetali, E. Vigoda, and L. Yang. 2013. Improving Mixing Conditions on the Grid for Counting and Sampling Independent Sets. *Probability Theory and Related Fields* 156 (2013), 75–99.
- [11] T.C. Schelling. 1969. Models of Segregation. *The American Economic Review* 59, 2 (1969), 488–493.
- [12] T.C. Schelling. 1971. Dynamic models of segregation. *The Journal of Mathematical Sociology* 1, 2 (1971), 143–186.
- [13] D. Vinković and A. Kirman. 2006. A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences* 103, 51 (2006), 19261–19265.