Stochastic Algorithms for Programmable Matter

JOSHUA J. DAYMUDE, ANDREA W. RICHA – ARIZONA STATE UNIVERSITY
SARAH CANNON – UNIVERSITY OF CALIFORNIA, BERKELEY
DANA RANDALL – GEORGIA INSTITUTE OF TECHNOLOGY

DISCRETE MATH SEMINAR – APRIL 3, 2019

Inspirations & Applications





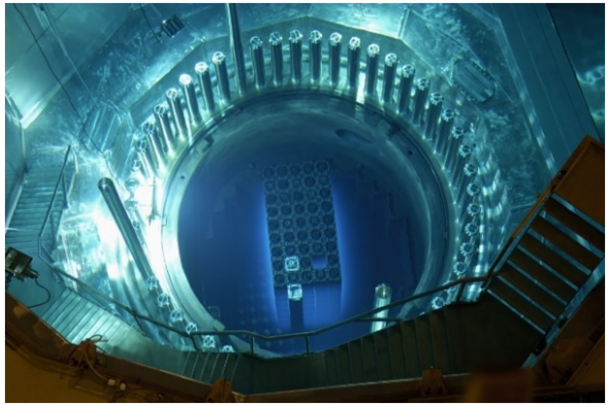
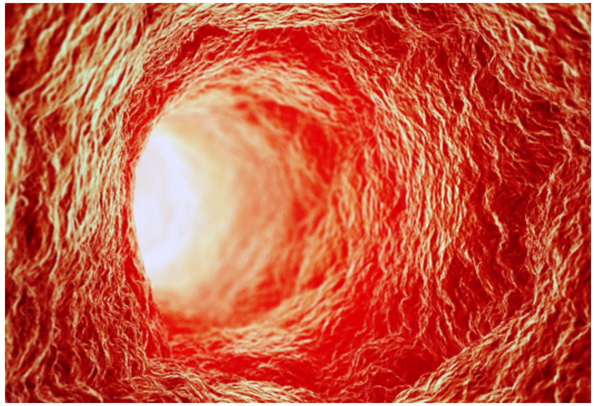

Current Programmable Matter

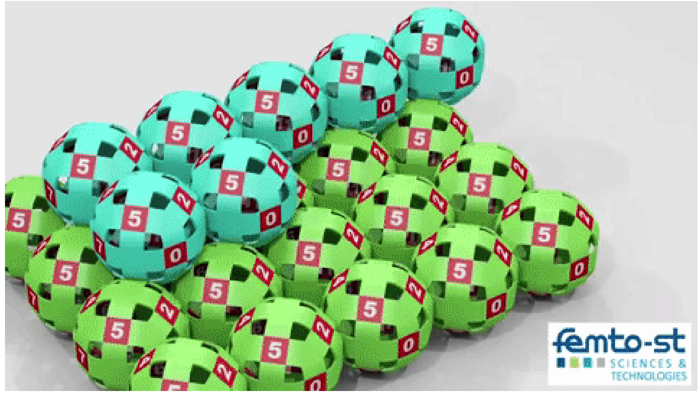

PB 2016: "Design of Quasi-Spherical Modules for Building Programmable Matter"

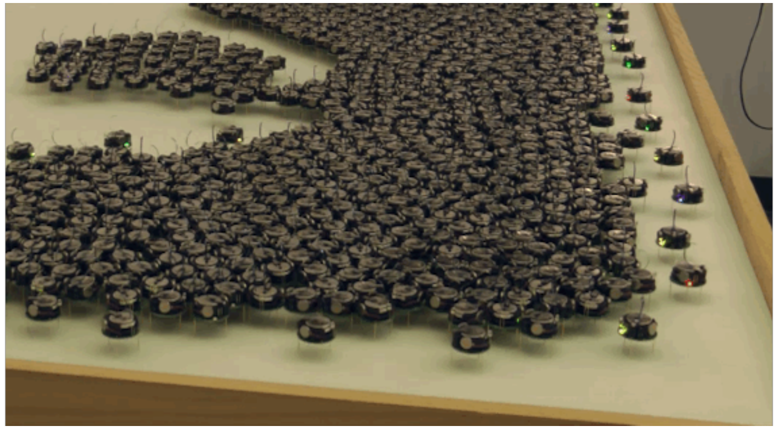

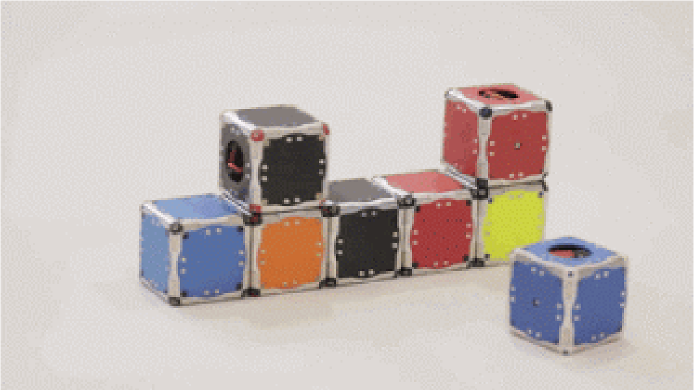

RGR 2013: "M-blocks: Momentum driven, magnetic modular robots"

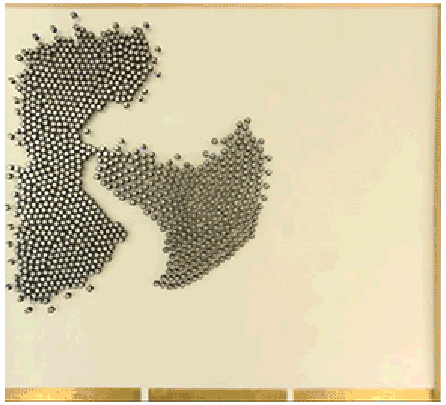

RCN 2014: "Programmable self-assembly in a thousand-robot swarm"

Current Programmable Matter

Programmable matter systems can be **passive** or **active**:

- **Passive**: Little/no control over decisions & movements, depends on the environment.
- **Active**: Can control actions & movements to solve problems.

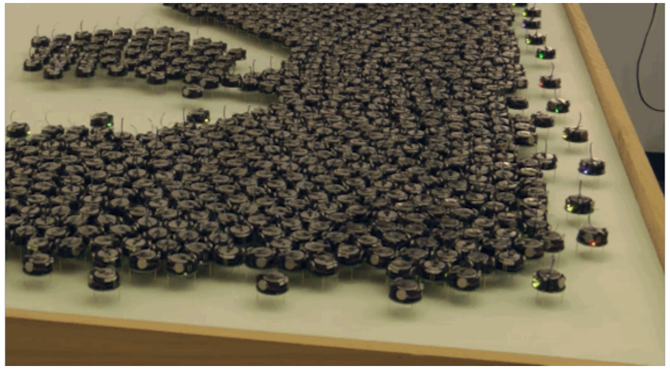

RCN 2014: "Programmable self-assembly in a thousand-robot swarm"

"Self-Organizing Particle Systems" (SOPS):

- Abstraction of **active** programmable matter.
- Each "particle" is a simple unit that can move and compute.
- Using **distributed algorithms**, limited particles coordinate to achieve sophisticated behavior.

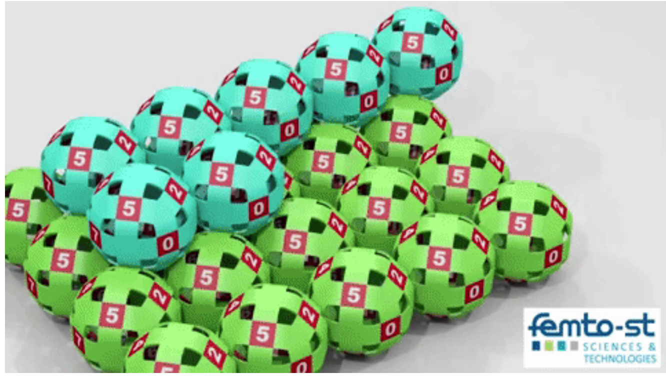

PB 2016: "Design of Quasi-Spherical Modules for Building Programmable Matter"

The Big Picture

What complex, collective behaviors are achievable by systems of simple, restricted programmable particles?

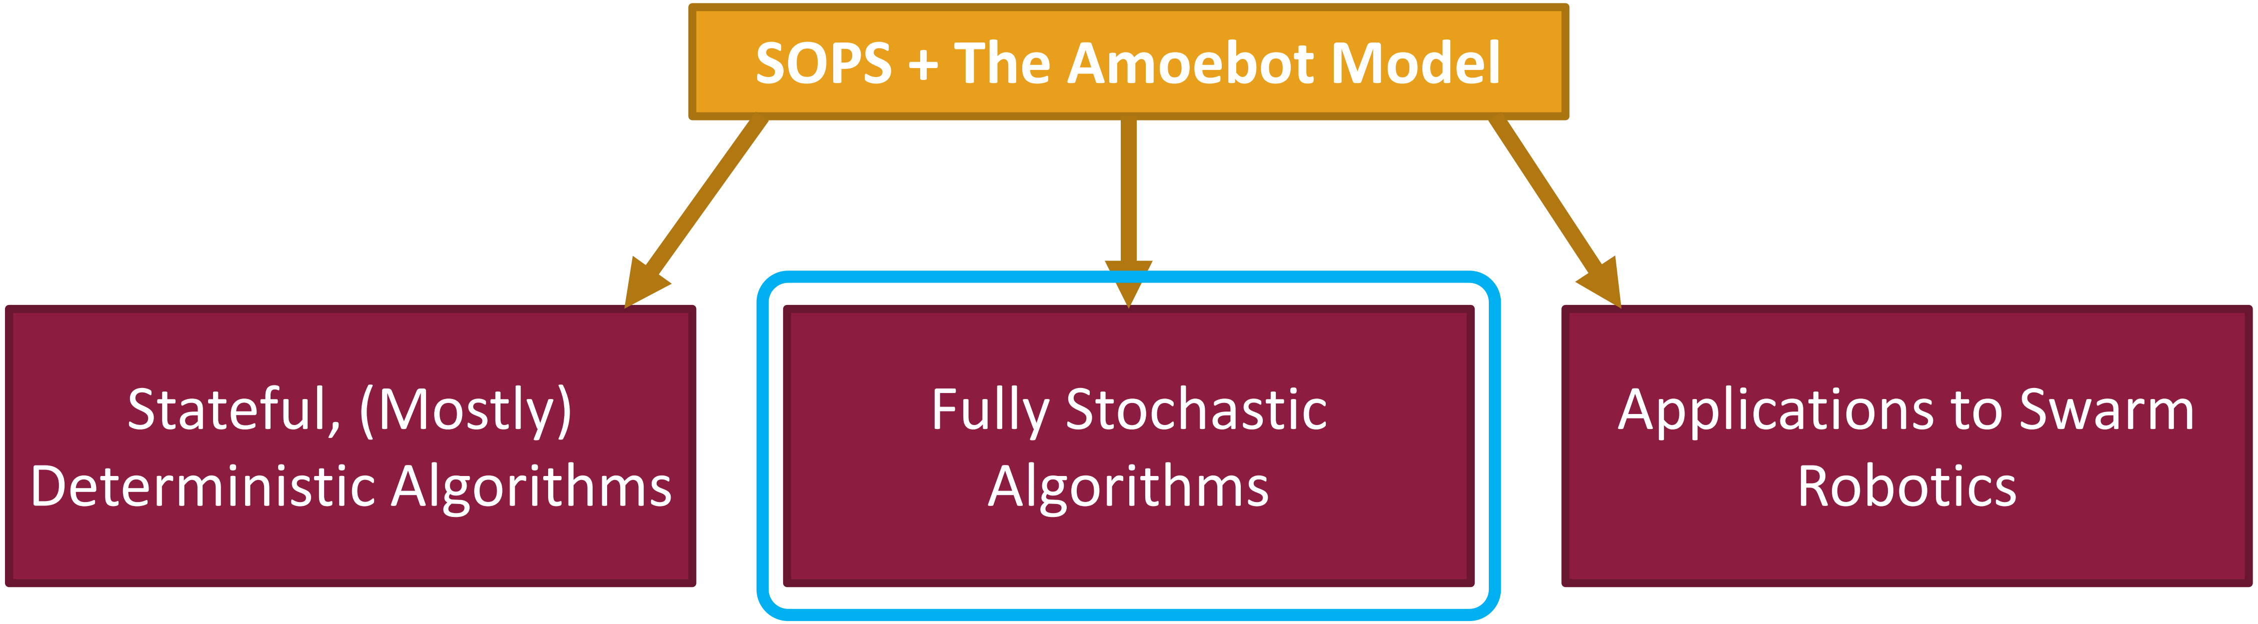

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.

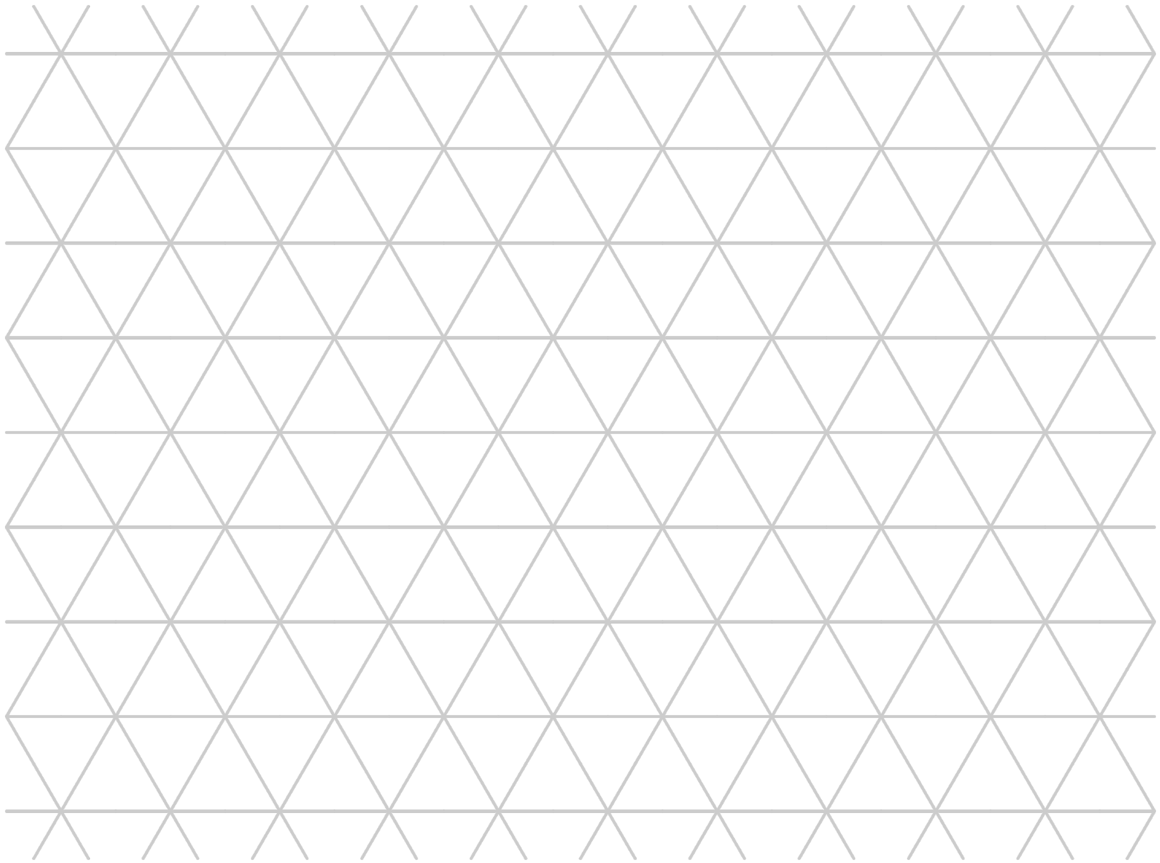

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted)...

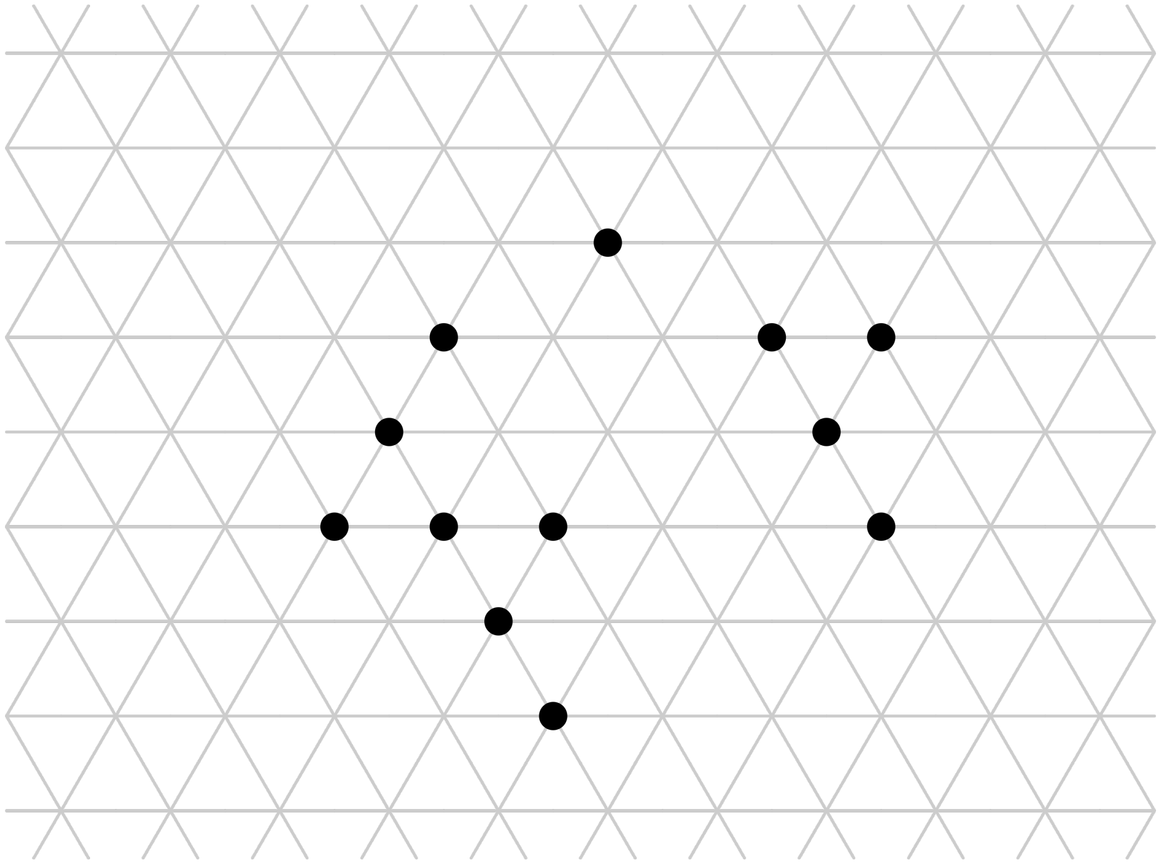

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted) or two adjacent nodes (expanded).

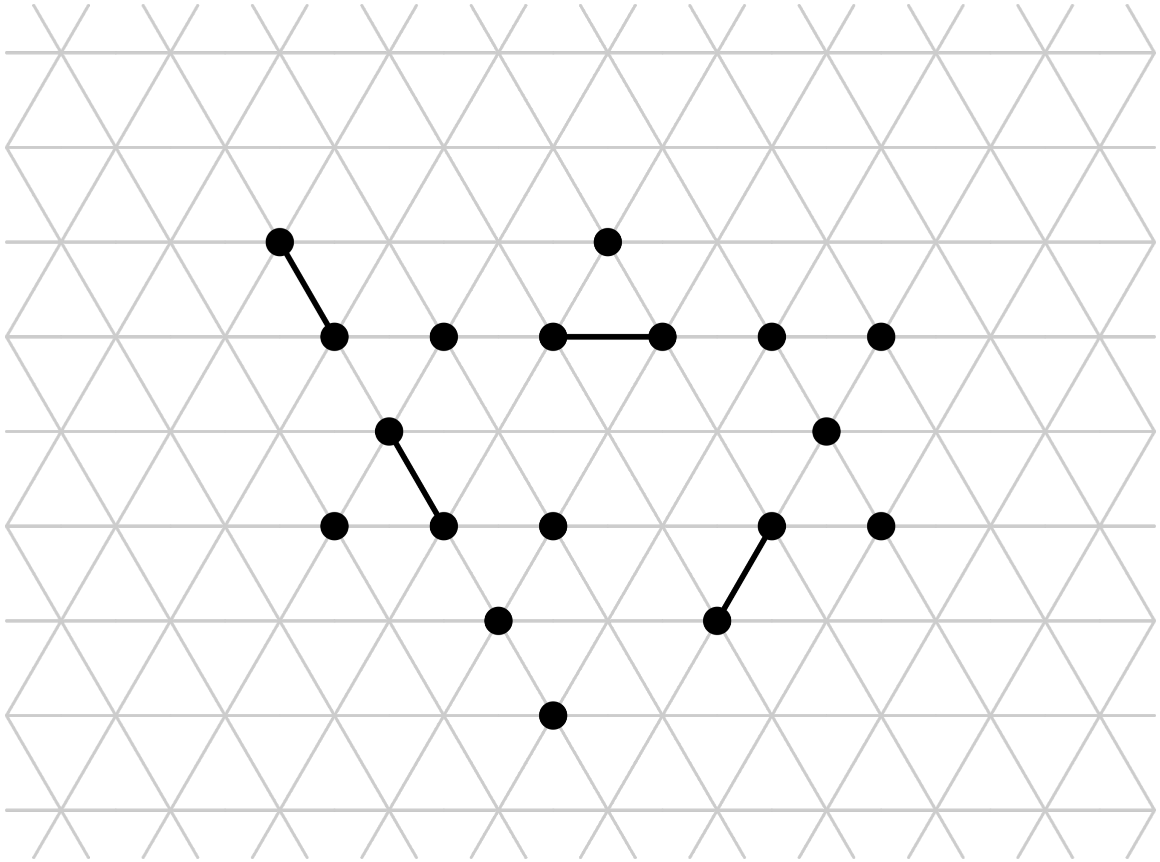

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted) or two adjacent nodes (expanded).
- Particles move by expanding and contracting.

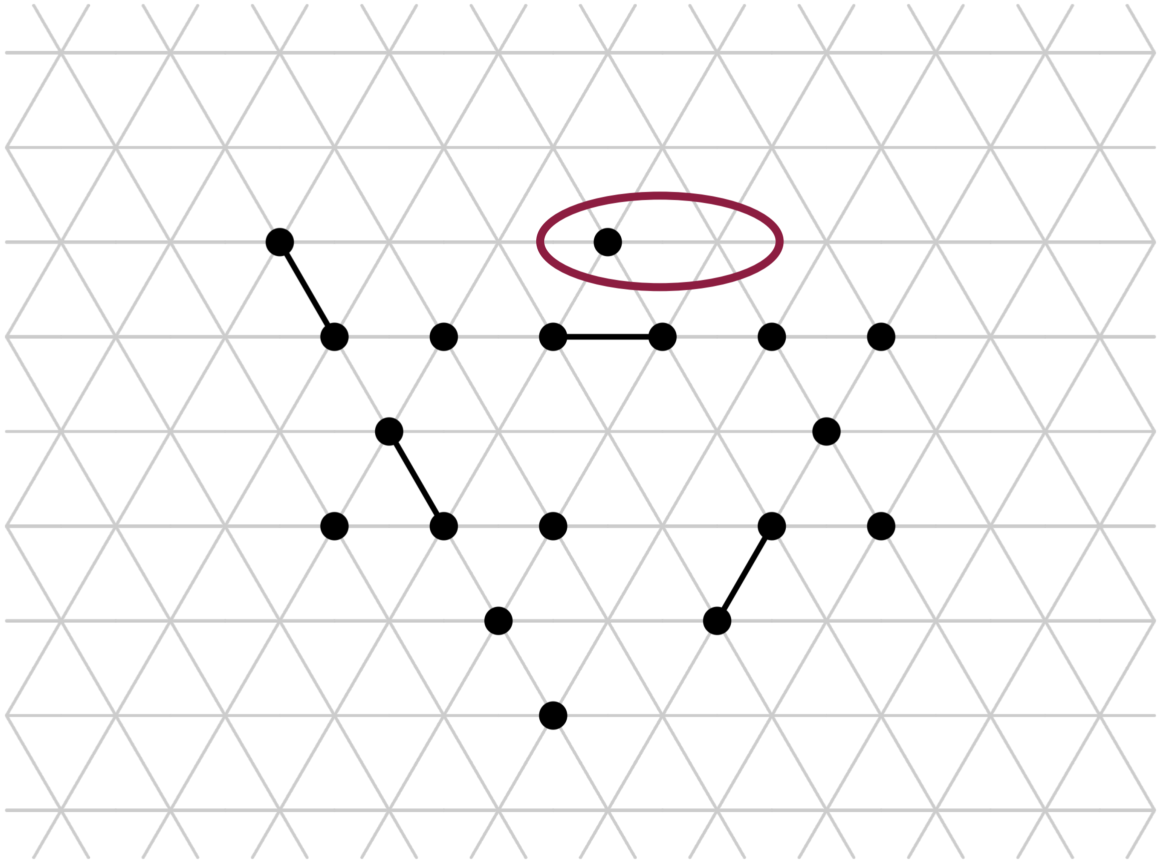

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted) or two adjacent nodes (expanded).
- Particles move by expanding and contracting.

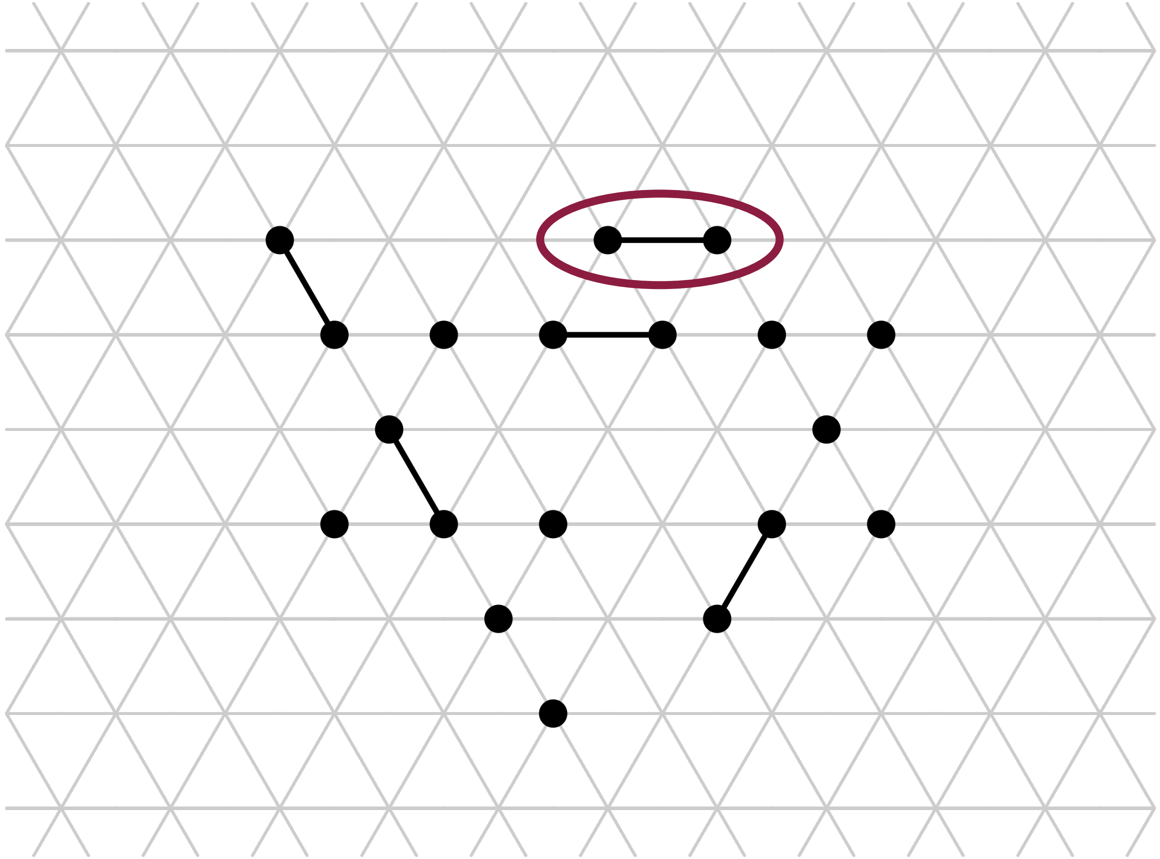

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted) or two adjacent nodes (expanded).
- Particles move by expanding and contracting.

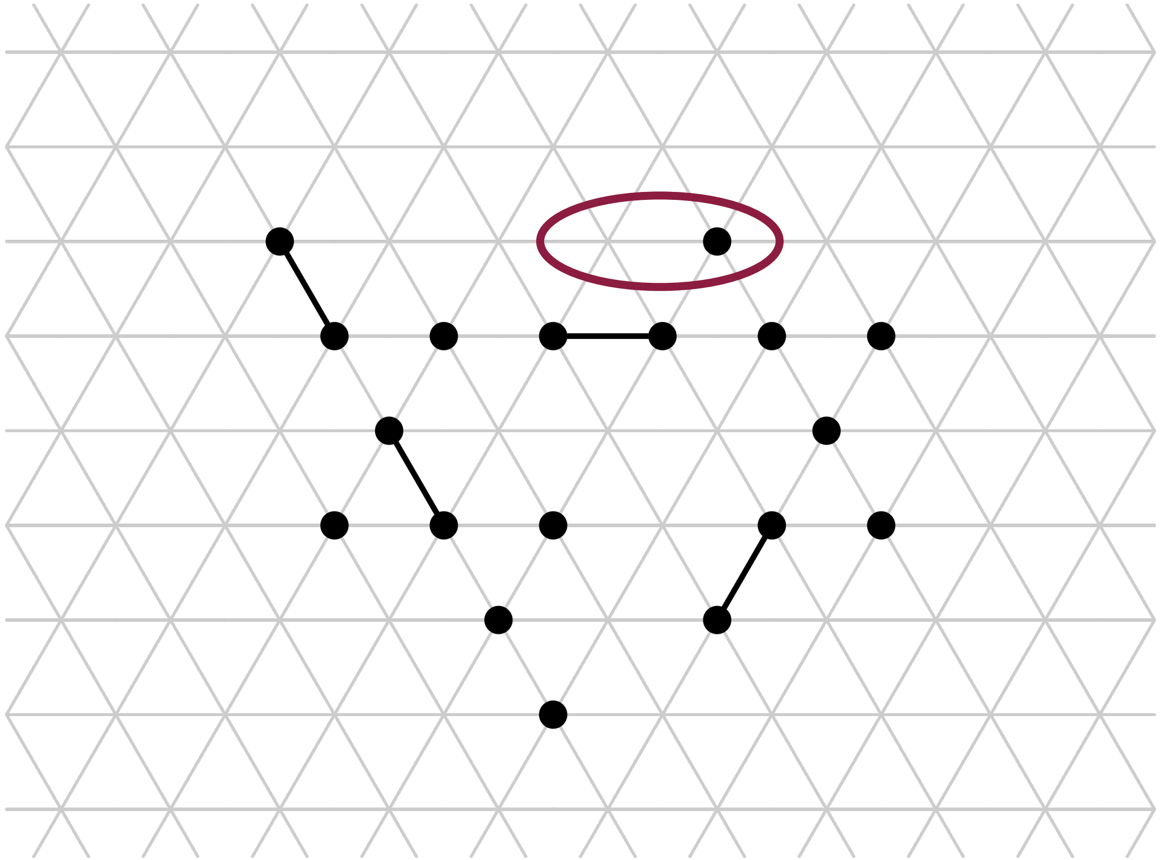

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted) or two adjacent nodes (expanded).
- Particles move by expanding and contracting.
- Particles do not have a global compass, but locally label their neighbors in clockwise order.

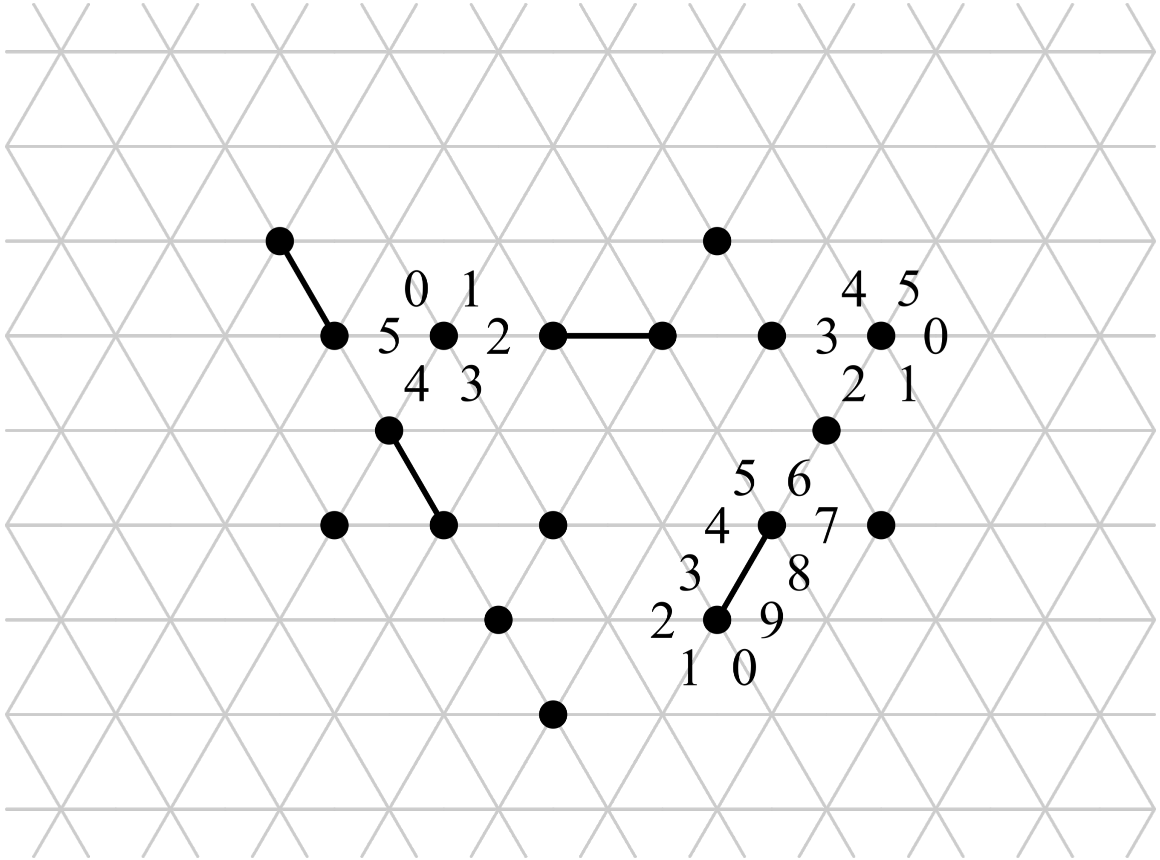

The (Geometric) Amoebot Model

- Space is modeled as the triangular lattice.
- Particles can occupy one node (contracted) or two adjacent nodes (expanded).
- Particles move by expanding and contracting.
- Particles do not have a global compass, but locally label their neighbors in clockwise order.
- Particles can communicate only with their neighbors.

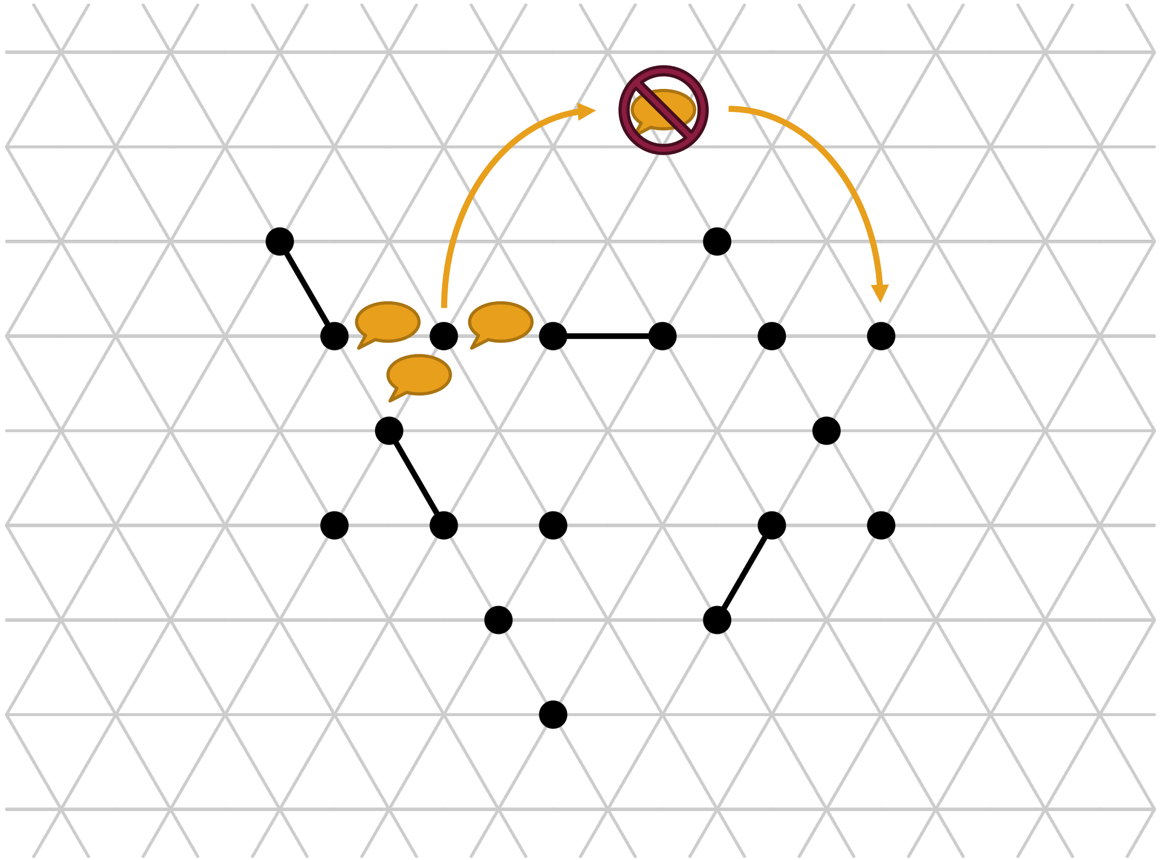

The (Geometric) Amoebot Model

- A particle only has constant-size memory.

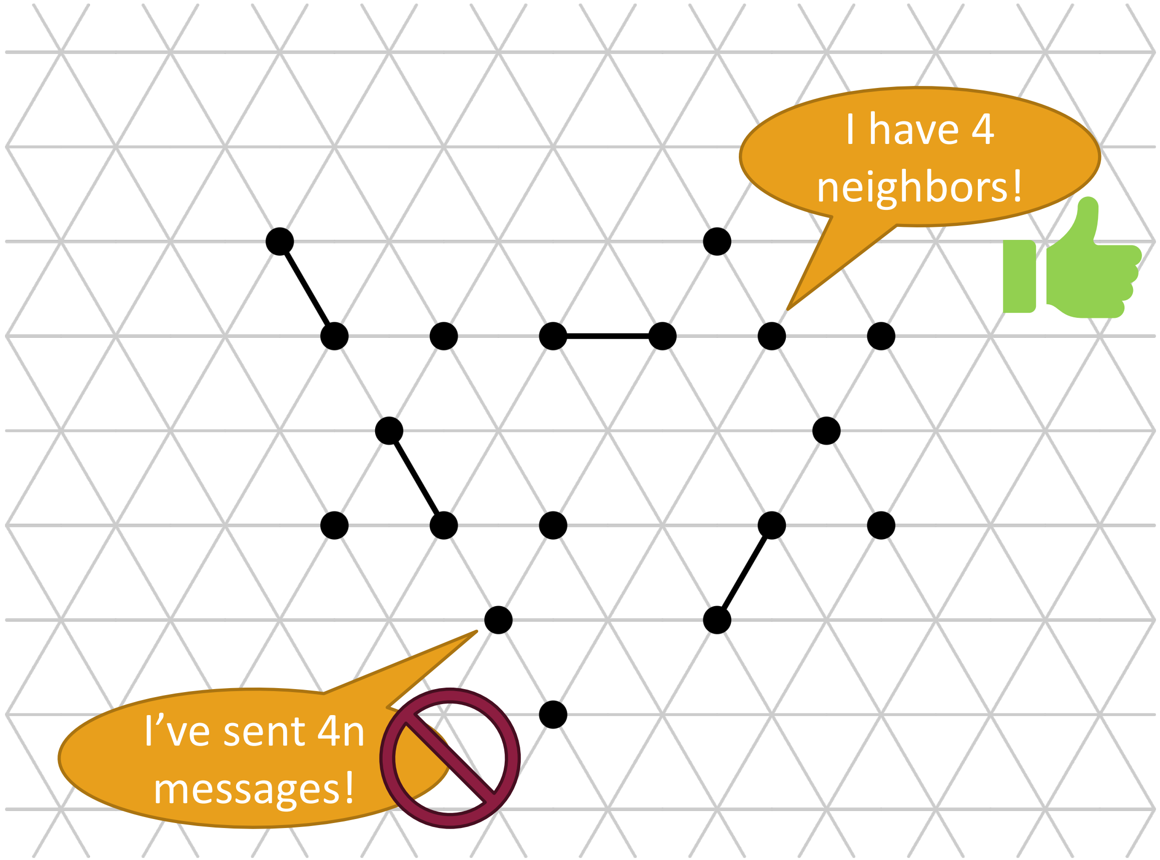

The (Geometric) Amoebot Model

- A particle only has constant-size memory.
- No unique identifiers.

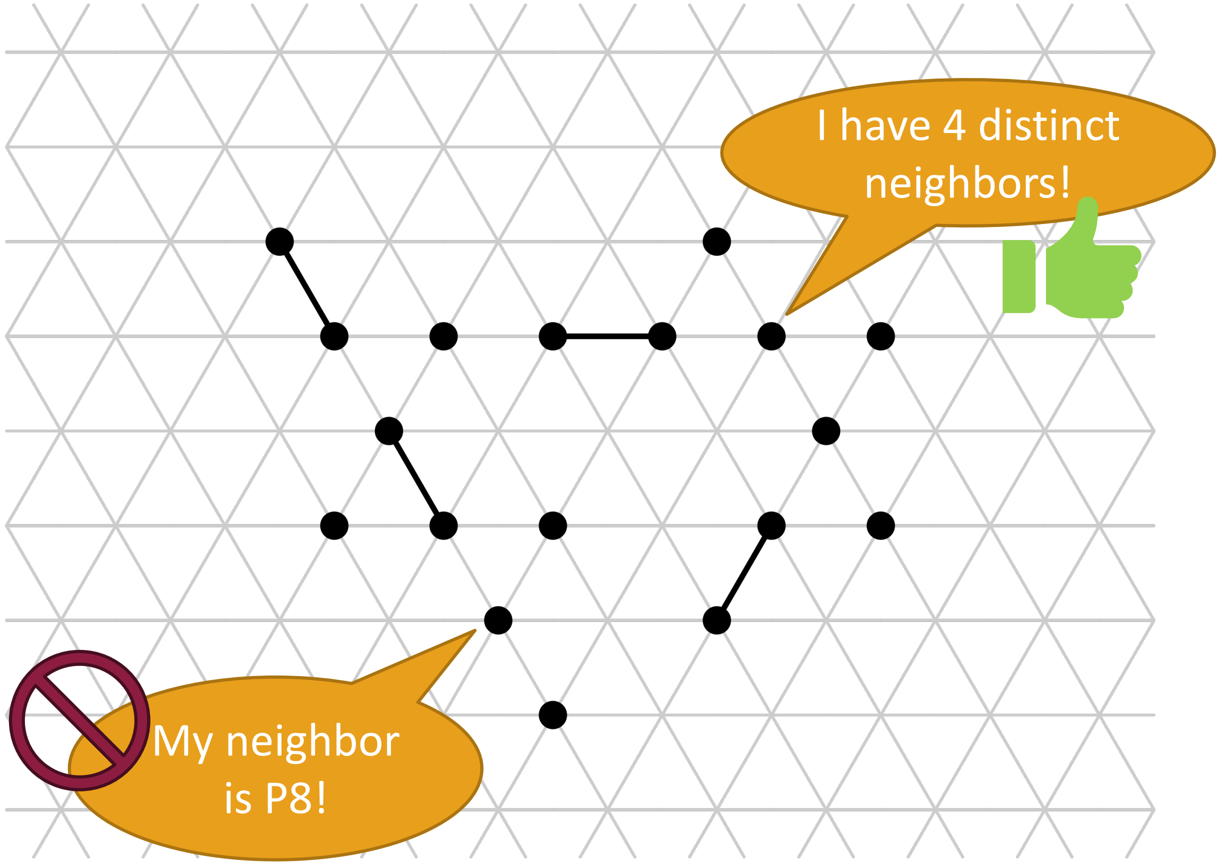

The (Geometric) Amoebot Model

- A particle only has constant-size memory.
- No unique identifiers.
- No global information.

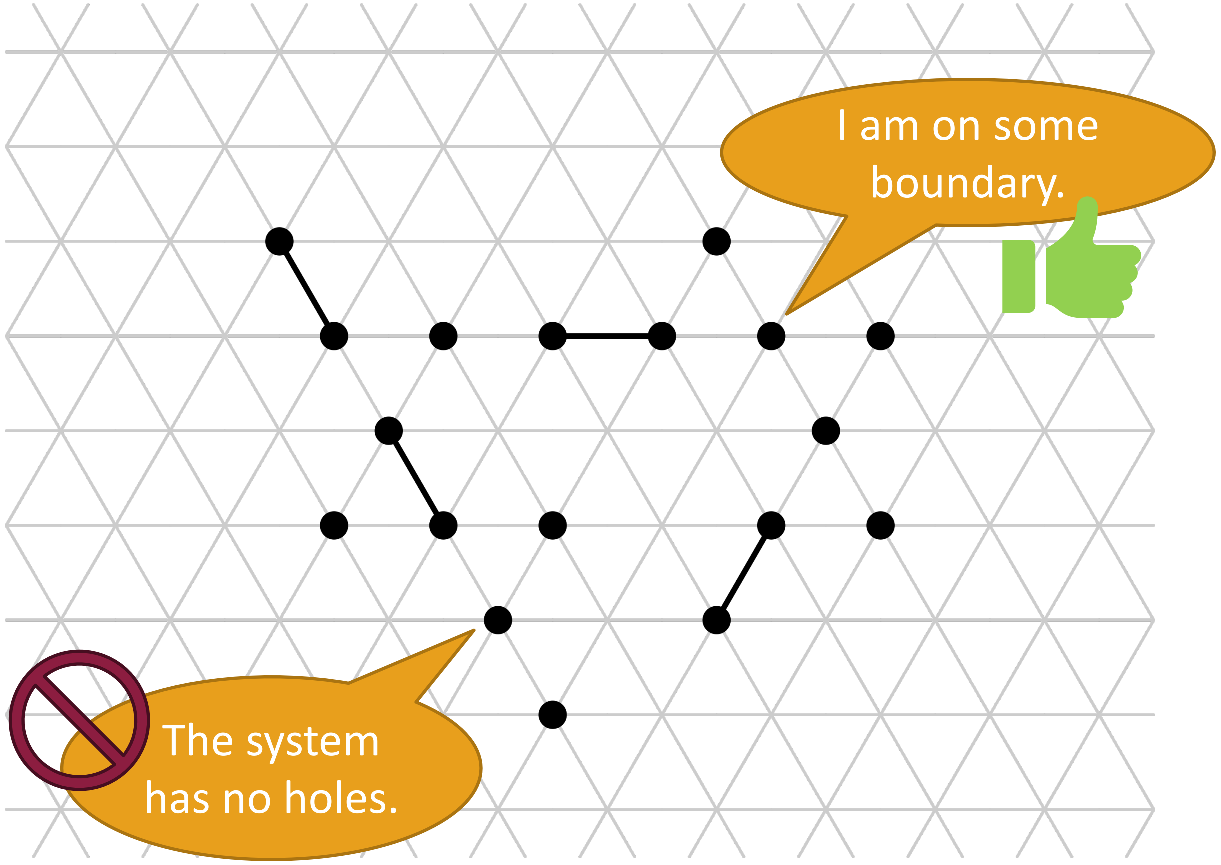

The (Geometric) Amoebot Model

- A particle only has constant-size memory.
- No unique identifiers.
- No global information.
- Asynchronous model of time: one atomic action may include finite computation and communication and at most one movement.

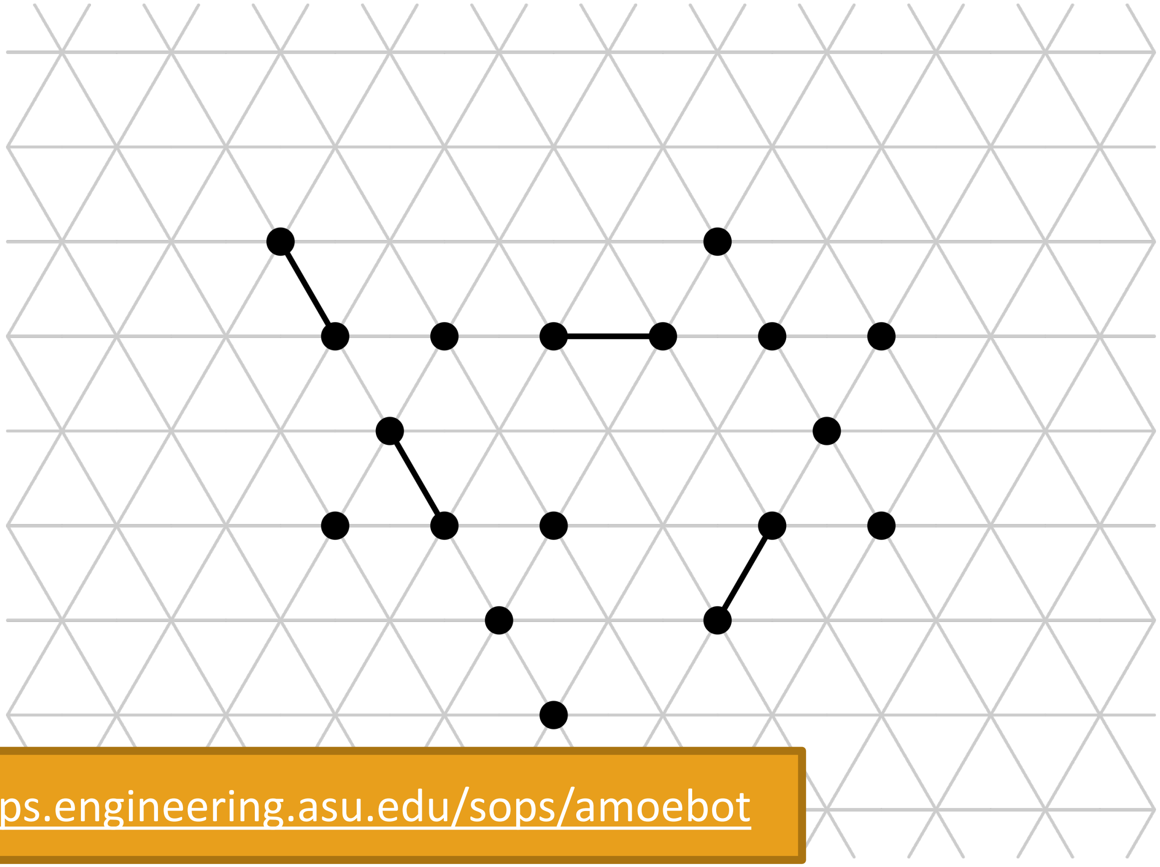

Read more at: sops.engineering.asu.edu/sops/amoebot

Markov Chains

- A Markov chain is a **memoryless**, **random** process that undergoes transitions between states in a state space.

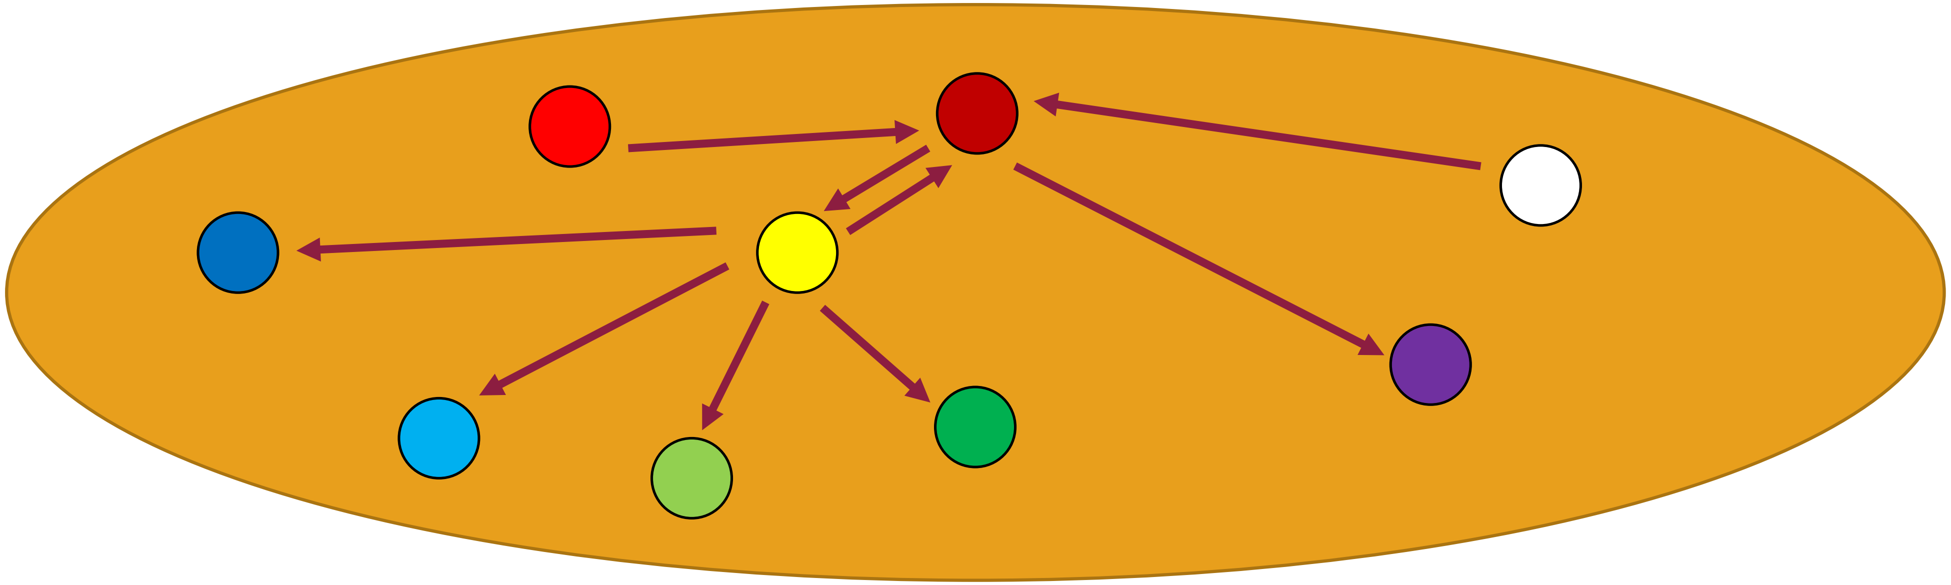

Markov Chains

- A Markov chain is a **memoryless**, **random** process that undergoes transitions between states in a state space.
- Our state space is all possible particle system **configurations**, and transitions between these configurations are individual **particle moves**.

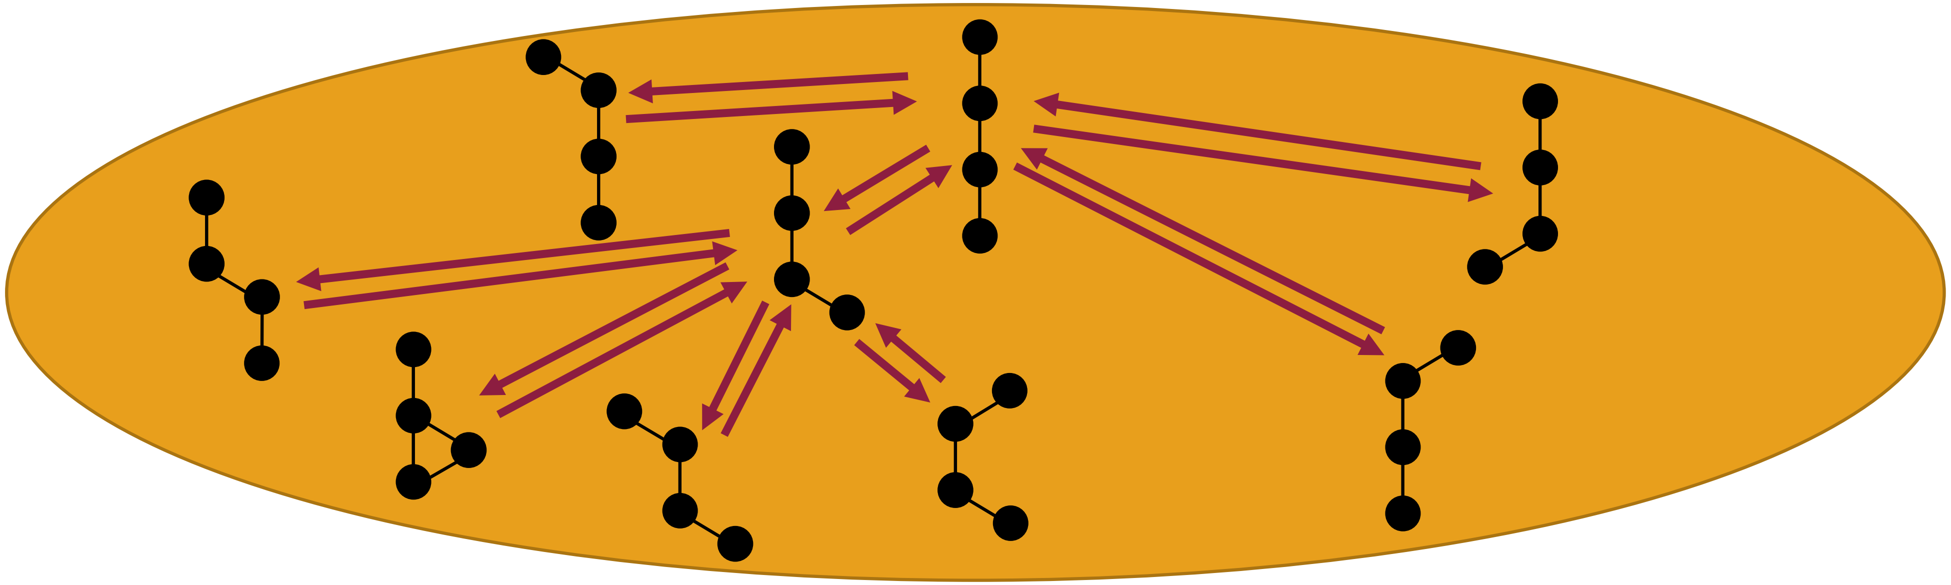

The Compression Problem

Informally: Gather a particle system P as tightly together as possible.

Formally:

- The **perimeter** of a connected, hole-free configuration σ, denoted p(σ), is the length of σ's outer boundary. Let $p_{min}(n)$ denote the minimum possible perimeter for n particles.
- Given a constant α > 1, σ is said to be **α-compressed** if $p(\sigma) \leq \alpha \cdot p_{min}(n)$.

"2-compressed"
$p(\sigma) = 23$
$p_{min}(19) = 12$

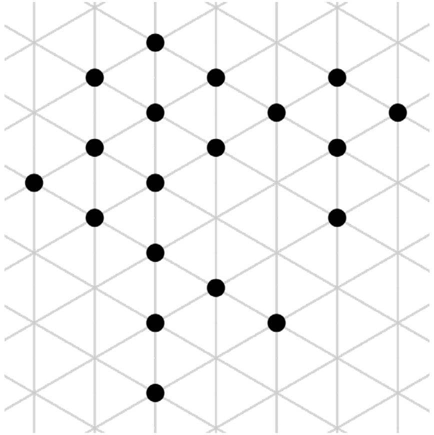

"1-compressed"
$p(\sigma) = 12$
$p_{min}(19) = 12$

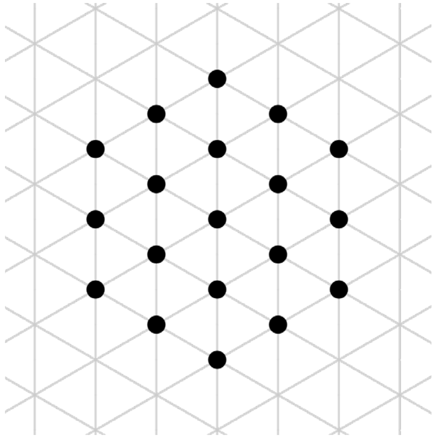

The Compression Problem

Given a particle system P of n particles in an arbitrary, connected configuration and a constant α > 1, reach and remain in a set of α-compressed configurations.

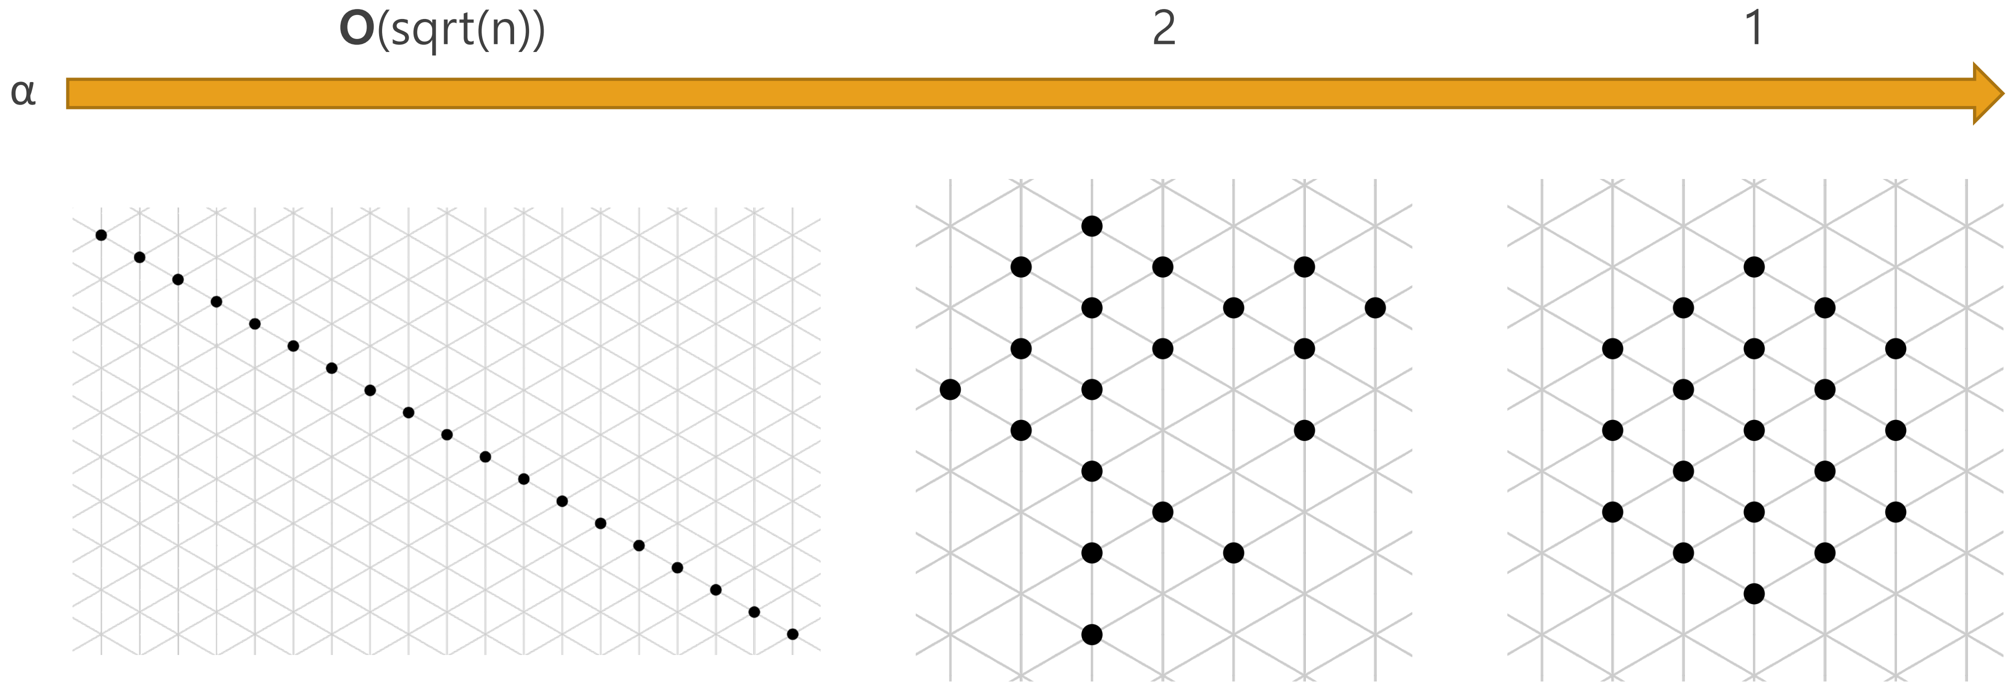

Why Stochastic?

Perimeter is a **global** property, but our particles are limited to **local** communication.

- **Lemma:** Maximizing the number of internal edges is equivalent to minimizing perimeter.
- First attempt: particles move to positions where they have more neighbors.

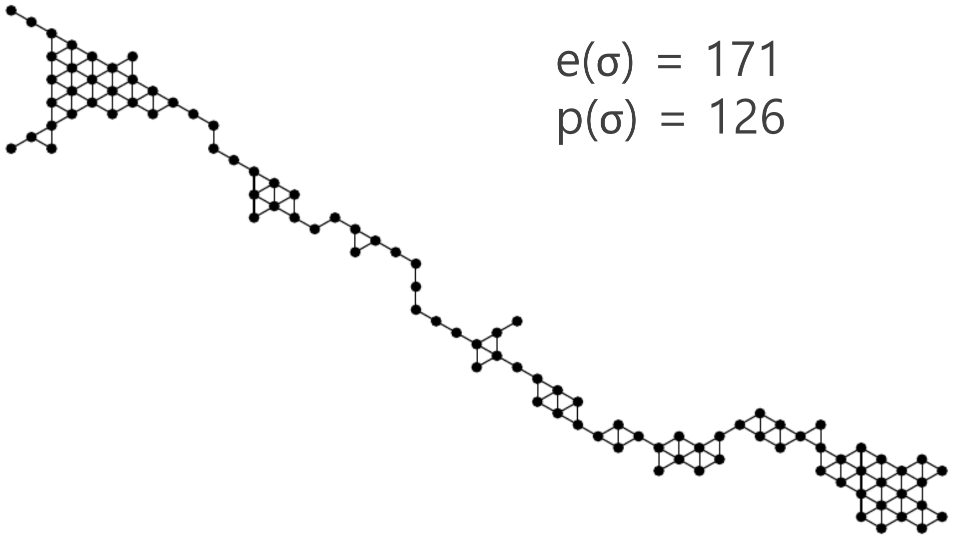

$e(\sigma) = 171$
$p(\sigma) = 126$

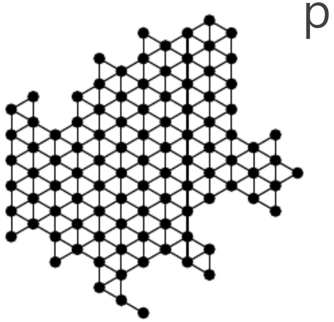

$e(\sigma) = 246$
$p(\sigma) = 51$

Why Stochastic?

Perimeter is a **global** property, but our particles are limited to **local** communication.

- **Lemma:** Maximizing the number of internal edges is equivalent to minimizing perimeter.
- First attempt: particles move to positions where they have more neighbors.
- ...However, we need something more robust to local minima.

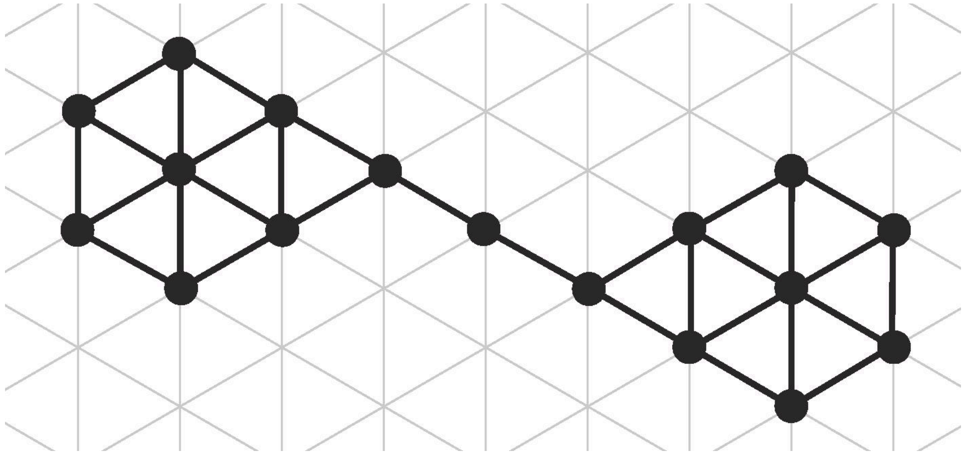

Markov Chains for Particle Systems

Turn a Markov chain (global, step-by-step) into a **local**, distributed, asynchronous algorithm:

- Carefully define the Markov chain to only use **local** moves.

Markov chain algorithm:

Starting from any configuration, repeat:

1. **Choose a particle** at random.
2. Expand into a (random) unoccupied adjacent position.
3. Perform some arbitrary, bounded computation involving its neighborhood.
4. Contract to either the new position or the original position.

Distributed algorithm:

Each particle concurrently executes:

2. Expand into a (random) unoccupied adjacent position.
3. Perform some arbitrary, bounded computation involving its neighborhood.
4. Contract to either the new position or the original position.

Markov Chains for Particle Systems

Turn a Markov chain (global, step-by-step) into a **local**, distributed, asynchronous algorithm:

- Carefully define the Markov chain to only use **local** moves.

Markov chain algorithm:

Starting from any configuration, repeat:

1. **Choose a particle** at random.
2. Expand into a (random) unoccupied adjacent position.
3. If certain properties hold, contract to new position with probability Pr[move].
4. Else, contract back to the original position.

Distributed algorithm:

Each particle concurrently executes:

2. Expand into a (random) unoccupied adjacent position.
3. If certain properties hold, contract to new position with probability Pr[move].
4. Else, contract back to the original position.

The Compression Algorithm

Input: an initial (connected) configuration σ_0 and a bias parameter $\boldsymbol{\lambda} > 1$.

Repeat:

1. Choose a particle from the system uniformly at random.
2. Choose an adjacent position uniformly at random. If occupied, go back to Step 1.
3. If properties hold for maintaining connectivity and avoiding holes, move to the chosen position with probability $\min\{1, \boldsymbol{\lambda}^{\Delta e}\}$.

Metropolis filter
(calculated w/ local info)

Proof: Detailed Balance

Theorem: We reach a stationary distribution π over configurations σ where $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)} \sim \boldsymbol{\lambda}^{e(\sigma)}$.

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

Proof: Peierls argument

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\boldsymbol{\lambda}^{-\Delta p}$) as a local change in #edges ($\boldsymbol{\lambda}^{\Delta e}$)!

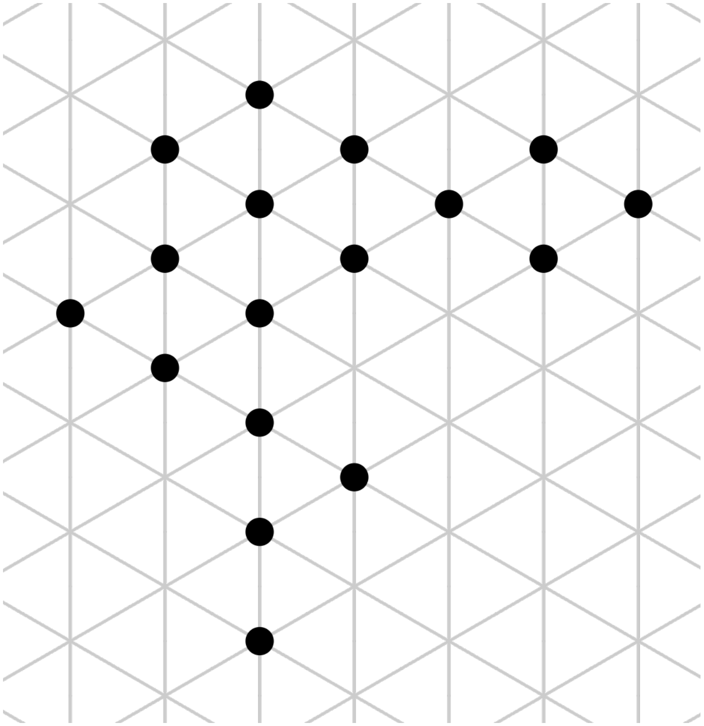

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\lambda^{-\Delta p}$) as a local change in #edges ($\lambda^{\Delta e}$)!

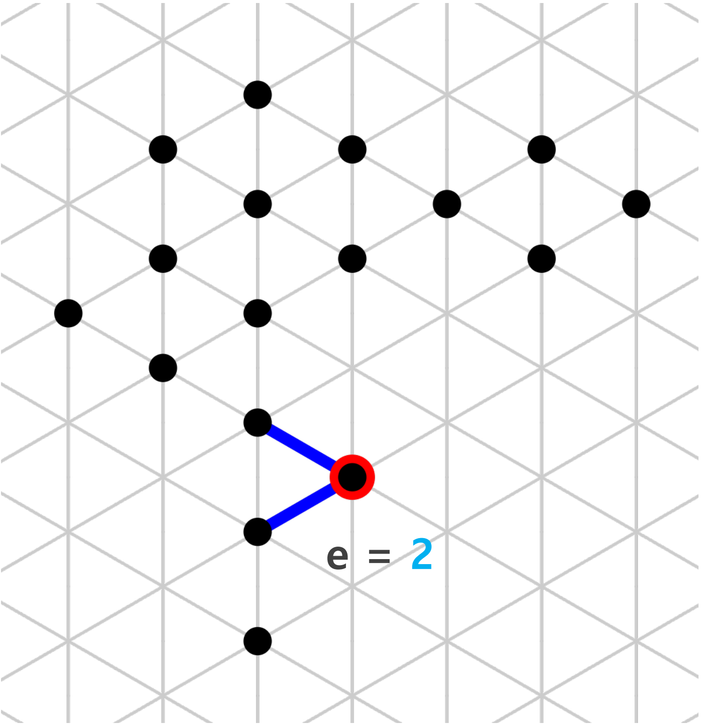

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\lambda^{-\Delta p}$) as a local change in #edges ($\lambda^{\Delta e}$)!

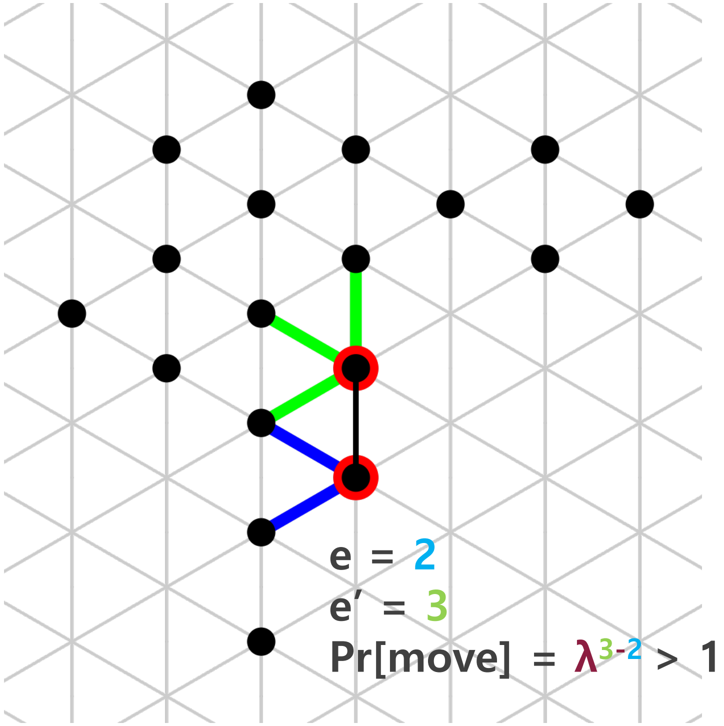

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\lambda^{-\Delta p}$) as a local change in #edges ($\lambda^{\Delta e}$)!

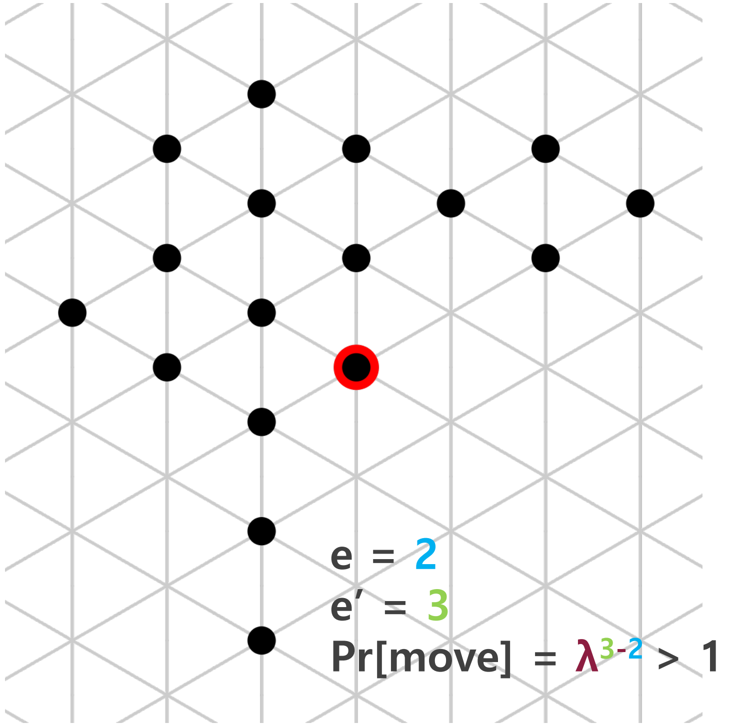

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\lambda^{-\Delta p}$) as a local change in #edges ($\lambda^{\Delta e}$)!

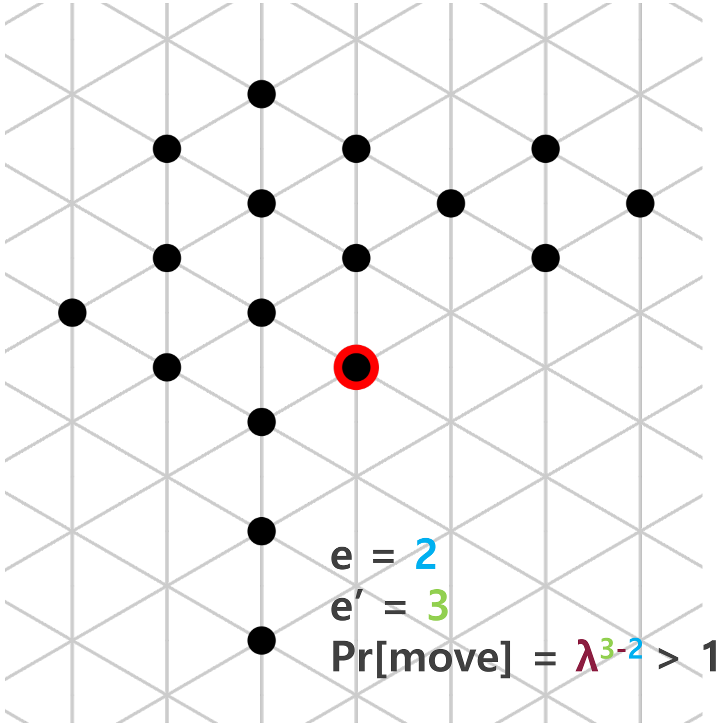

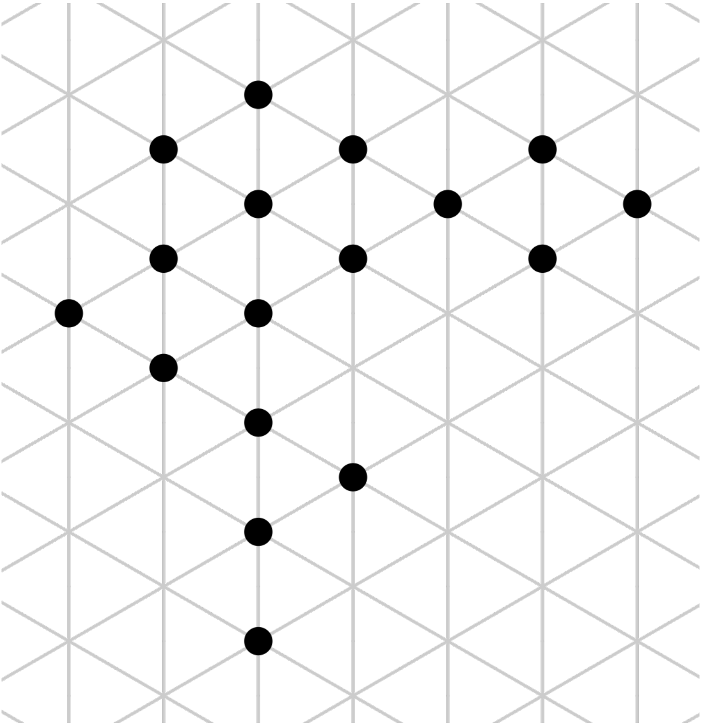

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\lambda^{-\Delta p}$) as a local change in #edges ($\lambda^{\Delta e}$)!

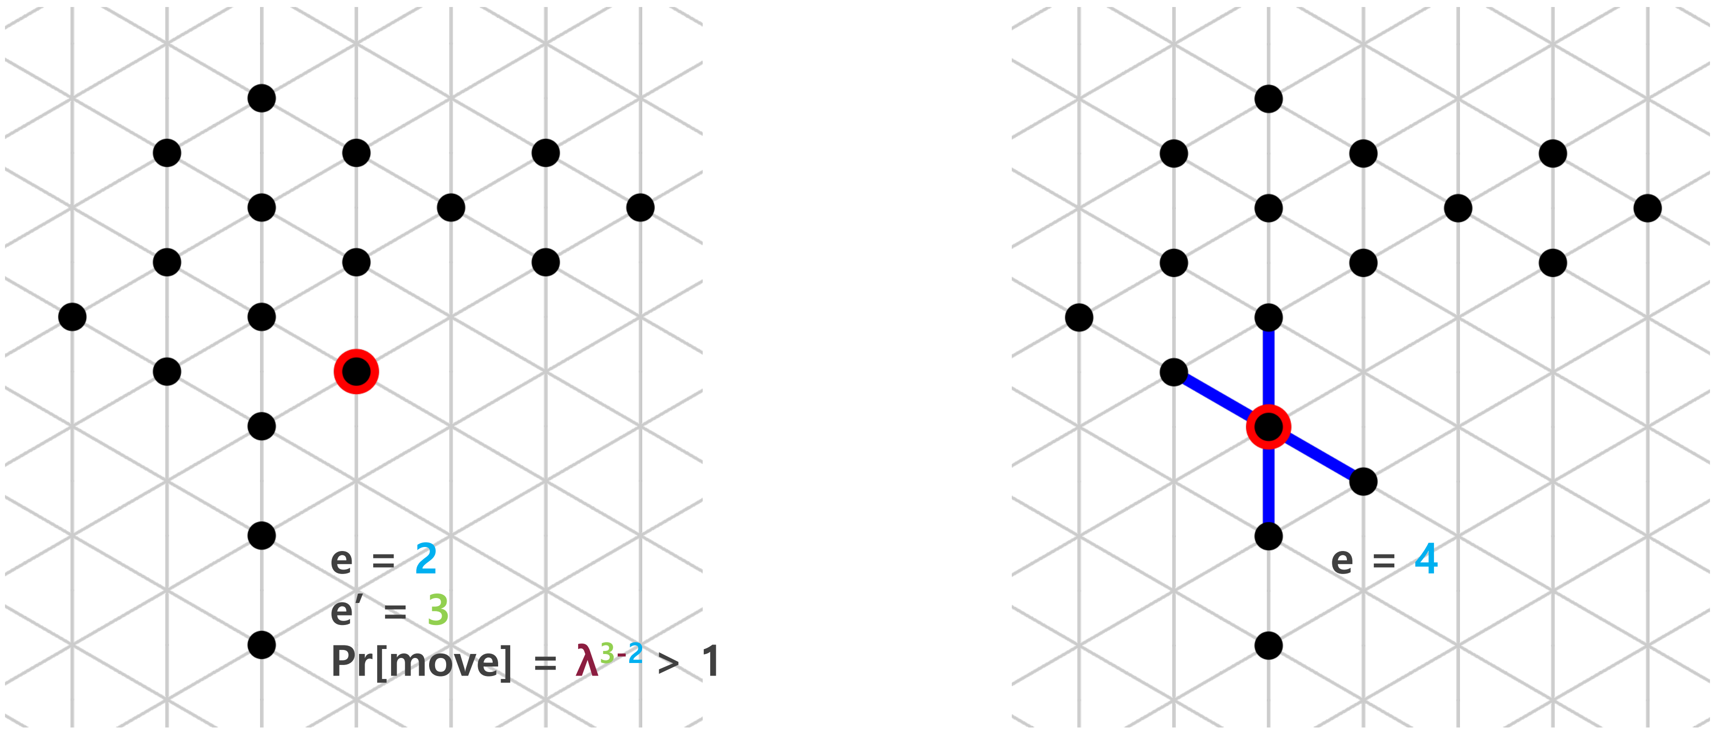

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, $e(\sigma) = 3n - p(\sigma) - 3$.

So, we can treat the global change in perimeter ($\boldsymbol{\lambda}^{-\Delta p}$) as a local change in #edges ($\boldsymbol{\lambda}^{\Delta e}$)!

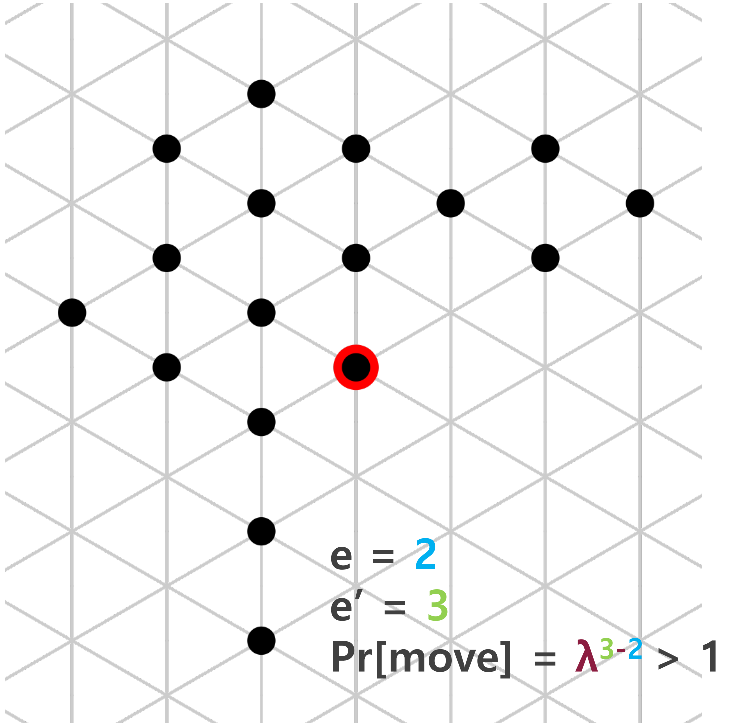

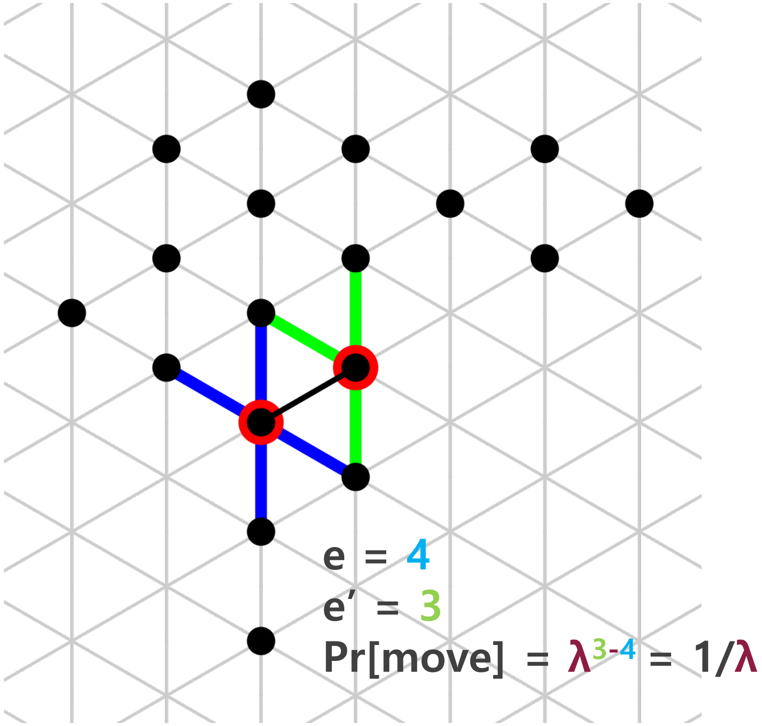

Compression: Example Runs

Lemma: For a connected, hole-free configuration σ of n particles, e(σ) = 3n – p(σ) – 3.

So, we can treat the global change in perimeter ($\lambda^{-\Delta p}$) as a local change in #edges ($\lambda^{\Delta e}$)!

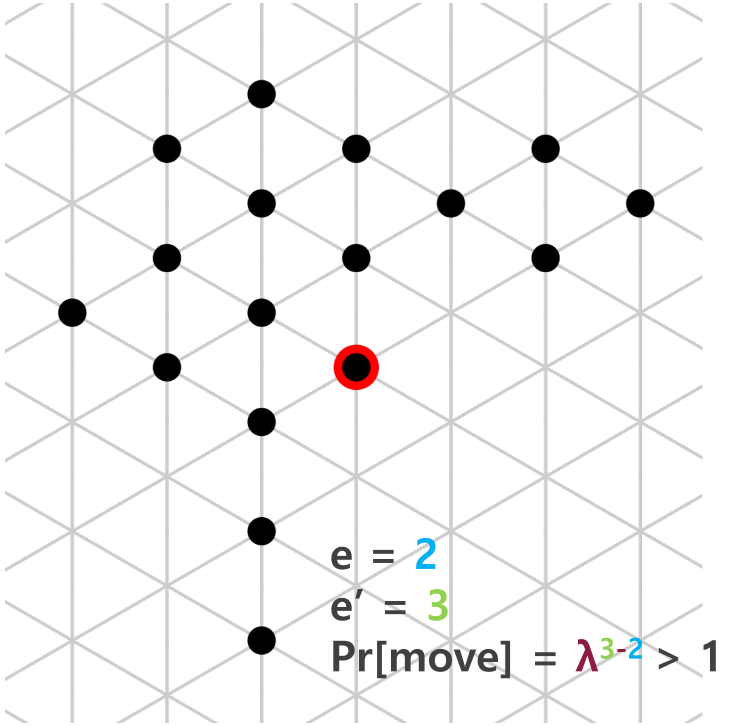

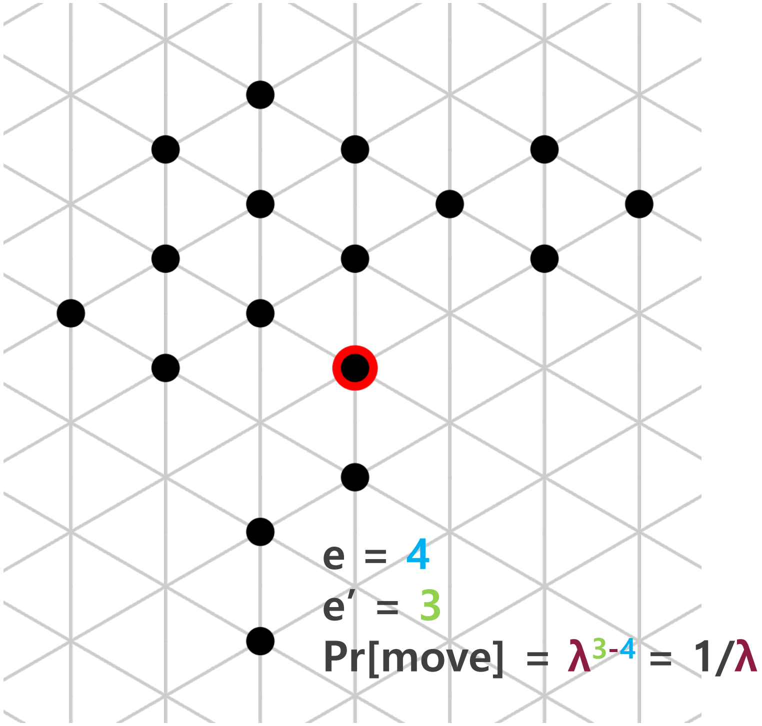

Compression: $\lambda = 4$

100 particles initially in a **line** after (a) 1 million, (b) 2 million, (c) 3 million, (d) 4 million, and (e) 5 million iterations.

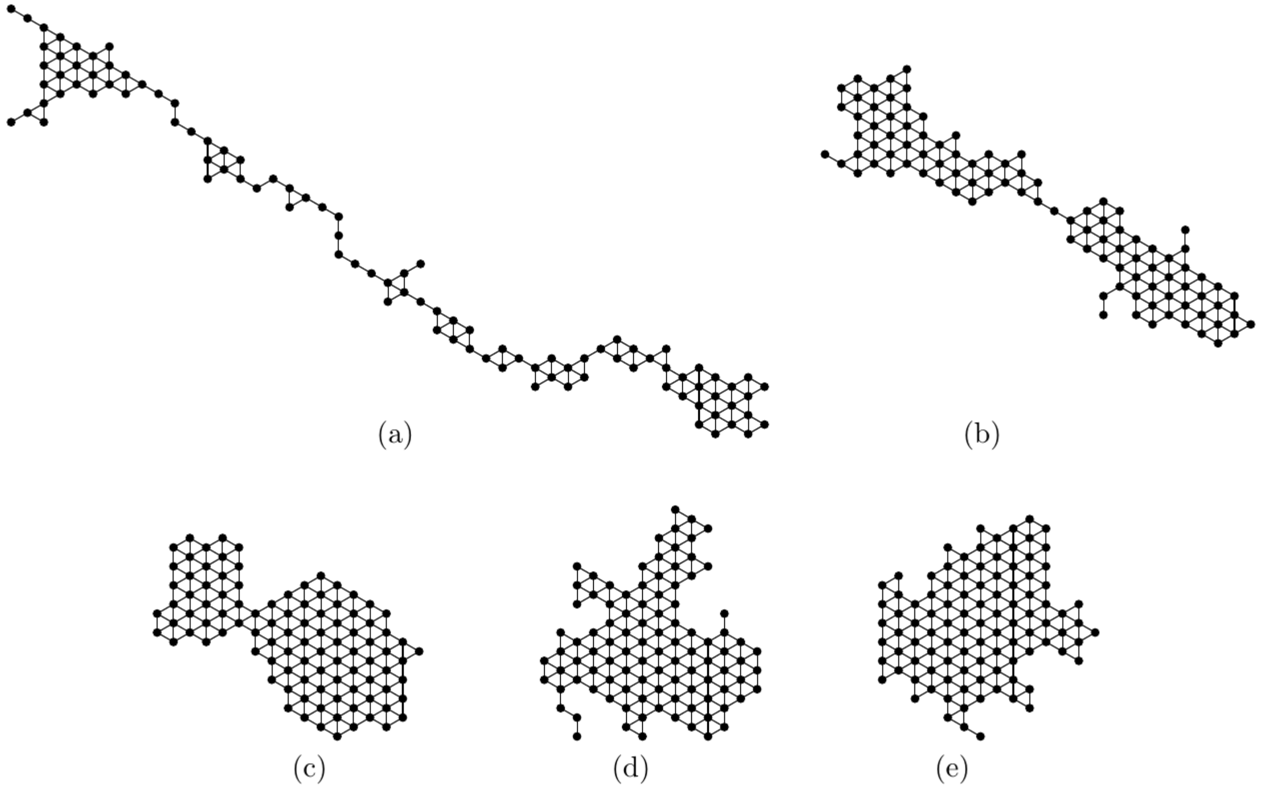

Local Properties for Movement

Input: an initial (connected) configuration σ_0 and a bias parameter $\lambda > 1$.

Repeat:

1. Choose a particle from the system uniformly at random.
2. Choose an adjacent position uniformly at random. If occupied, go back to Step 1.
3. If properties hold for maintaining connectivity and avoiding holes, move to the chosen position with probability $\min\{1, \boldsymbol{\lambda}^{\Delta e}\}$.

Theorem: We reach a stationary distribution π over configurations σ where $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)} \sim \boldsymbol{\lambda}^{e(\sigma)}$.

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

Local Properties for Movement

Qualitatively, what do we **not** want to happen to our particle system?

- The particle system could become **disconnected**.
- A (new) **hole** could be formed in the particle system.
- A move could be made that **couldn’t be “undone”** (bad for reversibility).

Local Properties for Movement

Allowed: "Slides" (with 1-2 pivots) and "jumps" (0 pivots) avoid bad outcomes.

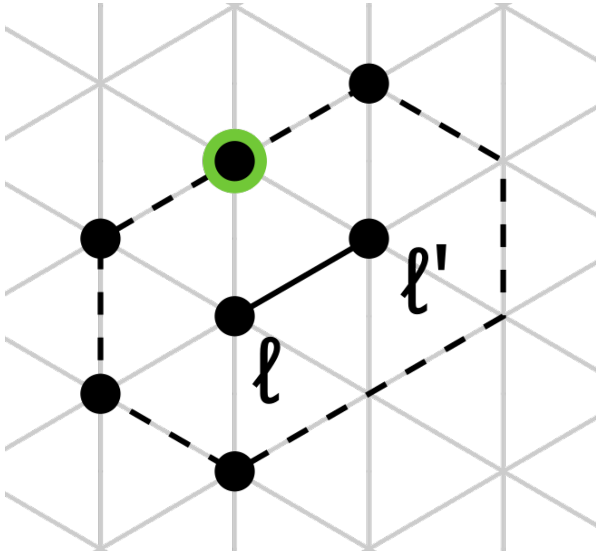

1 pivot

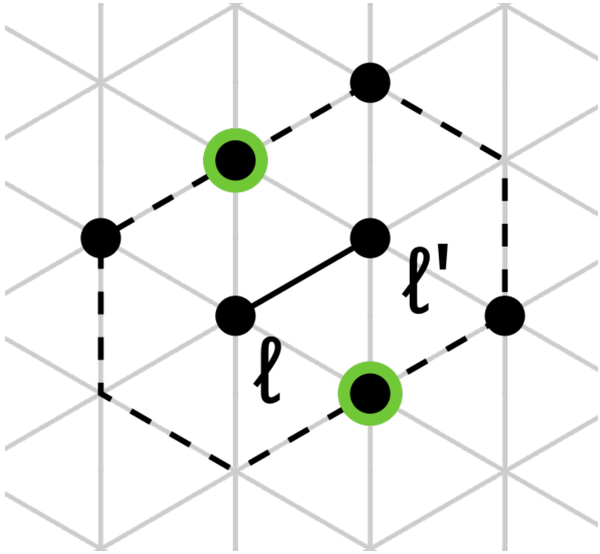

2 pivots

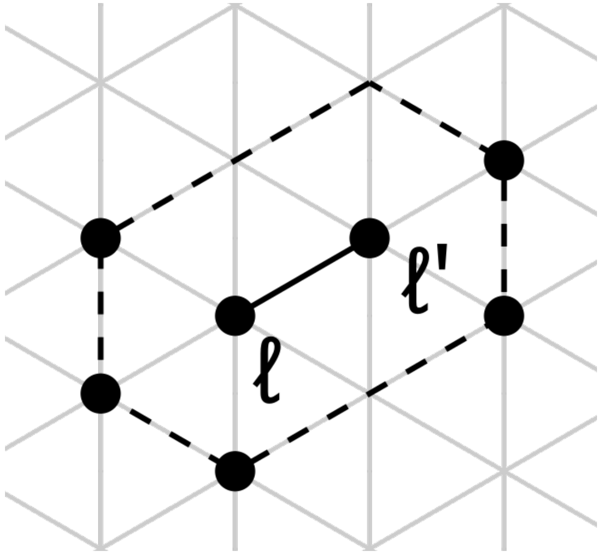

no pivot

Local Properties for Movement

Not allowed: Moves that lead to disconnection, create new holes, or cannot be reversed.

1. The current location should not have 5 neighbors (moving forms a hole).
2. If there are 1-2 pivots, all neighbors should be locally connected to a pivot.
3. If there are 0 pivots, both locations should have locally connected neighborhoods.

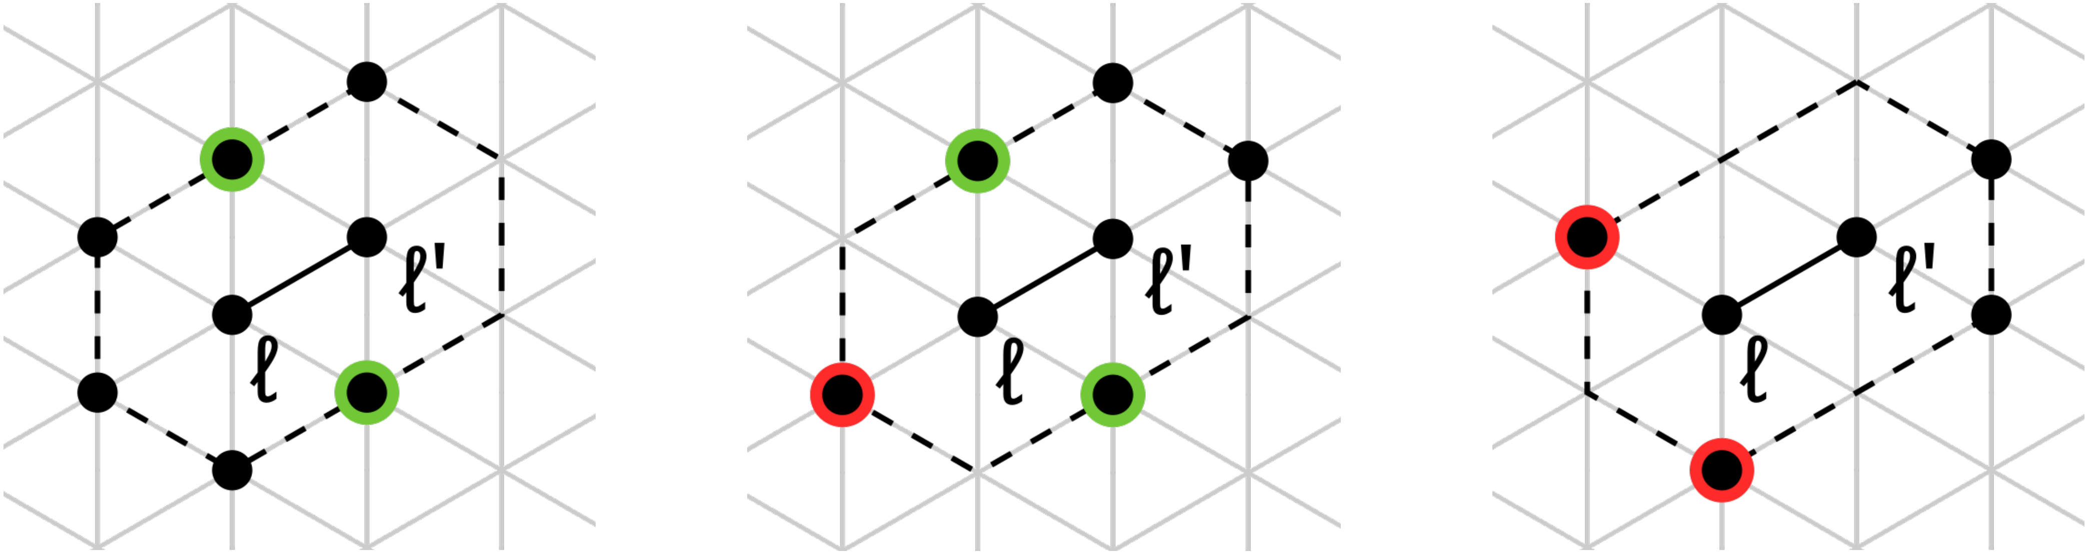

The Stationary Distribution

Input: an initial (connected) configuration σ_0 and a bias parameter $\lambda > 1$.

Repeat:

1. Choose a particle from the system uniformly at random.
2. Choose an adjacent position uniformly at random. If occupied, go back to Step 1.
3. If properties hold for maintaining connectivity and avoiding holes, move to the chosen position with probability $\min\{1, \boldsymbol{\lambda}^{\Delta e}\}$.

Theorem: We reach a stationary distribution π over configurations σ where $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)} \sim \boldsymbol{\lambda}^{e(\sigma)}$.

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

The Stationary Distribution

With the local rules for movement, our algorithm has the following properties:

- The particle system remains connected and no new holes form.
- **Lemma:** All existing holes are eventually eliminated.
- Once all holes are eliminated, all moves are reversible.

Irreducible + Aperiodic -> Ergodic -> Unique Stationary Distribution

Theorem: Our Markov chain for compression is ergodic on the state space of all connected, hole-free configurations.

Thus, our Markov chain for compression has a unique stationary distribution π.

The Stationary Distribution

Theorem: We reach a stationary distribution π over configurations σ where $\pi(\sigma) = \lambda^{e(\sigma)} / Z$.

"The Metropolis filter $\min\{1, \lambda^{\Delta e}\}$ actually gets us the stationary distribution we wanted."

Proof.

- π is the stationary distribution if $\pi(\sigma) \cdot P(\sigma,\tau) = \pi(\tau) \cdot P(\tau,\sigma)$, the **detailed balance condition**.
- Suppose, w.l.o.g., that $\lambda^{e(\tau) - e(\sigma)} \leq 1$. Then:

$$\begin{aligned}\pi(\sigma) \cdot P(\sigma,\tau) &= (\lambda^{e(\sigma)} / Z) \cdot (1/n) \cdot (1/6) \cdot \min\{ 1, \lambda^{e(\tau) - e(\sigma)} \} \\ &= (\lambda^{e(\sigma) + e(\tau) - e(\sigma)} / Z) \cdot (1/n) \cdot (1/6) \\ &= (\lambda^{e(\tau)} / Z) \cdot (1/n) \cdot (1/6) \cdot 1 \\ &= \pi(\tau) \cdot P(\tau,\sigma)\end{aligned}$$

Correctness: A Peierls Argument

Input: an initial (connected) configuration σ_0 and a bias parameter $\lambda > 1$.

Repeat:

1. Choose a particle from the system uniformly at random.
2. Choose an adjacent position uniformly at random. If occupied, go back to Step 1.
3. If properties hold for maintaining connectivity and avoiding holes, move to the chosen position with probability $\min\{1, \boldsymbol{\lambda}^{\Delta e}\}$.

Theorem: We reach a stationary distribution π over configurations σ where $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)} \sim \boldsymbol{\lambda}^{e(\sigma)}$.

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

Correctness: A Peierls Argument

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

"In the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$, the α-compressed configurations are most likely."

<u>Proof sketch.</u>

- Let S_α be the set of configurations σ with $p(\sigma) > \alpha \cdot p_{min}(n)$ (the bad ones).
- We show, at stationarity, it is exponentially unlikely to be in such a "bad" configuration:

$$\pi(S_\alpha) \leq d^{sqrt(n)}, \text{ where } d < 1.$$

- Let A_k be the set of "bad" configurations with $p(\sigma) = k$.
- The weight of a configuration σ in A_k is $\boldsymbol{\lambda}^{-k}$.
- But how many configurations are in A_k?

Correctness: A Peierls Argument

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

Proof sketch (cont.)

- How many configurations are in A_k?
- **Lemma:** There are at most $f(k)(2 + \mathrm{sqrt}(2))^k$ configurations in A_k, where f is subexponential.
- So we can calculate $\pi(A_k)$ as follows:

$$\pi(A_k) \;=\; \boldsymbol{\lambda}^{-k} \cdot |A_k| \,/\, Z \;\leq\; \boldsymbol{\lambda}^{-k} \cdot f(k)(2 + \mathrm{sqrt}(2))^k \,/\, Z$$

Correctness: A Peierls Argument

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

Proof sketch (cont.)

- We have $\pi(A_k) \leq \boldsymbol{\lambda}^{-k} \cdot f(k)(2 + \mathrm{sqrt}(2))^k / Z$.
- We can easily lower bound $Z = \sum_\sigma \boldsymbol{\lambda}^{-p(\sigma)} \geq \boldsymbol{\lambda}^{-pmin(n)}$.
- Finally, sum $\pi(A_k)$ over all perimeters k from $\alpha \cdot p_{min}(n)$ to $p_{max}(n) = 2n - 2$.
- Carrying out a lot of algebra, we get:

$$\pi(S_\alpha) = \sum_{k = \alpha \cdot pmin : 2n-2} \pi(A_k) \leq \sum_k \boldsymbol{\lambda}^{-k} \cdot f(k)(2 + \mathrm{sqrt}(2))^k / \boldsymbol{\lambda}^{-pmin(n)} \leq \ldots \leq d^{\mathrm{sqrt}(n)}, \text{ where } d < 1.$$

Correctness: A Peierls Argument

Theorem: For any $\alpha > 1$, there is a $\boldsymbol{\lambda}$ (depending on α) so that in the stationary distribution $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ we have $p(\sigma) \leq \alpha \cdot p_{min}(n)$ with high probability.

Corollary: For any $\boldsymbol{\lambda} > 2 + \text{sqrt}(2)$ with $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$, there is an $\alpha > 1$ such that, at stationarity, with all but exponentially small probability the particle system is α-compressed.

But surprisingly…

Theorem: For any $\boldsymbol{\lambda} < 2.17$ with $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ and any $\alpha > 1$, at stationarity, the probability that the particle system is α-compressed is exponentially small.

"Expanding" Beyond Compression

The last theorem shows that setting **λ** < 2.17 yields the opposite behavior: **expansion**.

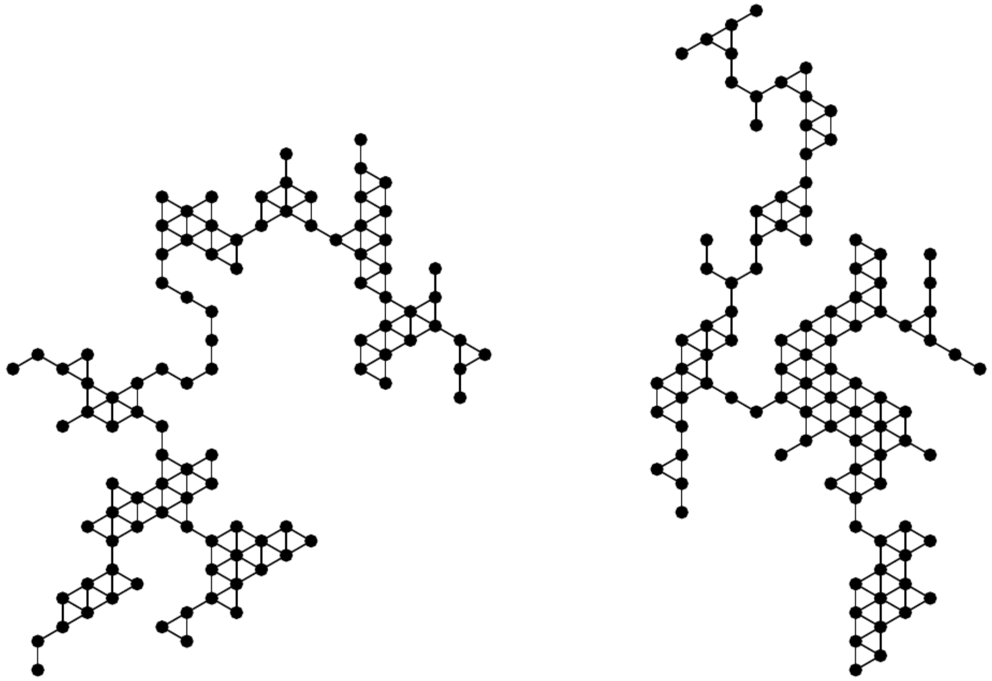

(a) (b)

What else can we do?

Shortcut Bridging

Problem Statement: Maintain bridge structures that simultaneously balance the tradeoff between the benefit of a shorter path and the cost of more particles in the bridge.

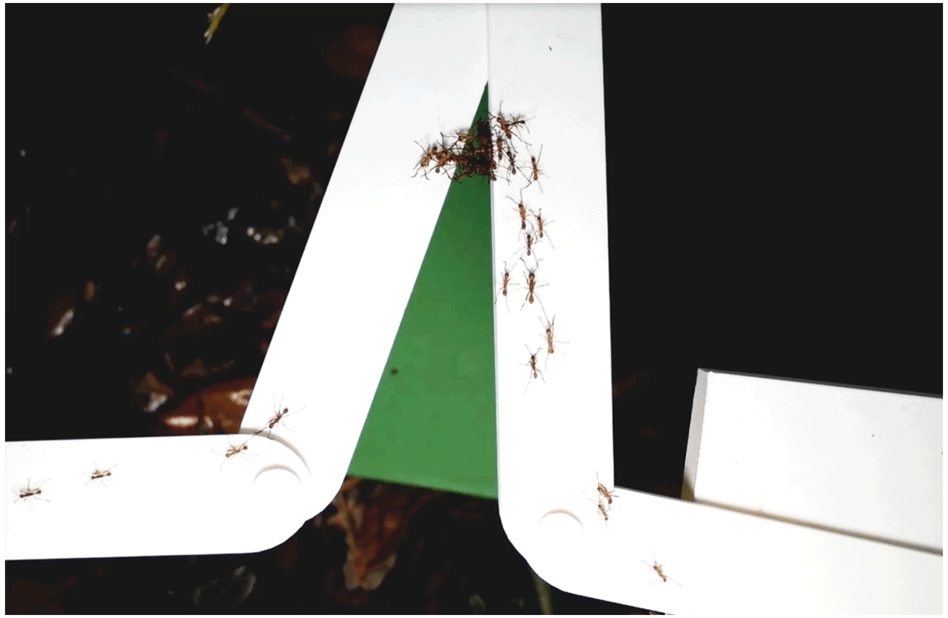

RLPKCG 2015: "Army ants dynamically adjust living bridges..."

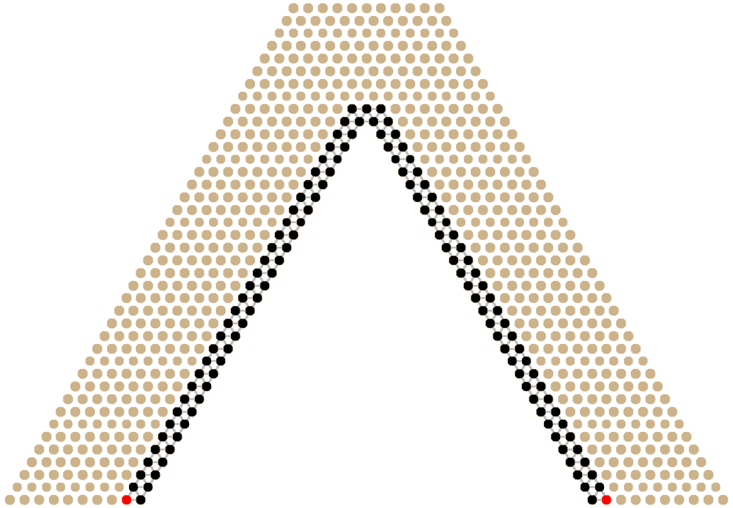

Shortcut Bridging

We achieve this goal by extending the compression algorithm to minimize both the total perimeter $p(\sigma)$ and the gap perimeter $g(\sigma)$.

Formally, we minimize **weighted perimeter** $p'(\sigma,c) = p(\sigma) + c \cdot g(\sigma)$, where $c > 0$.

Min $p(\sigma)$

Max $g(\sigma)$

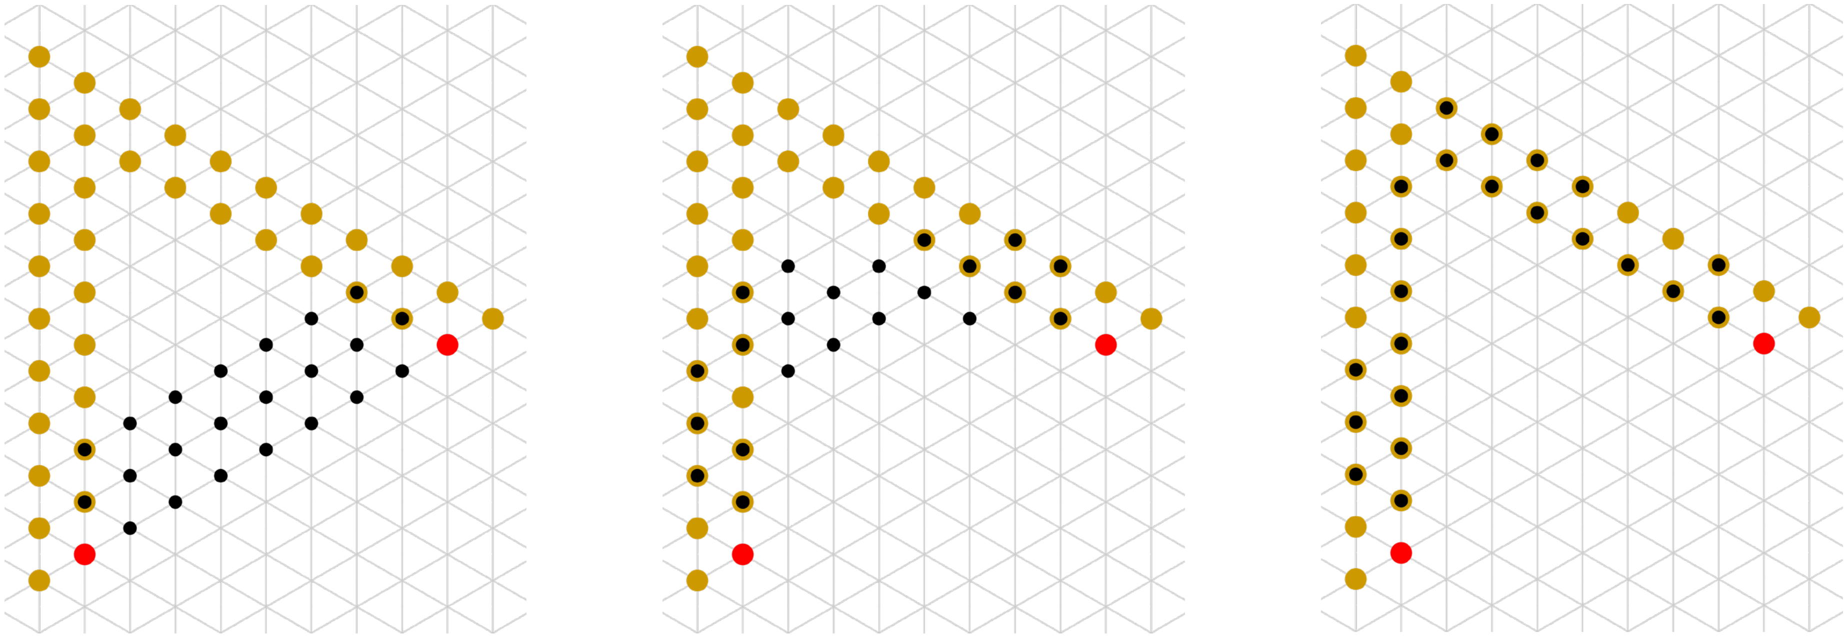

Max $p(\sigma)$

Min $g(\sigma)$

Shortcut Bridging: λ = 4, γ = 2

A particle system initially fully on a V-shaped land mass after (a) 2 million, (b) 4 million, (c) 6 million, and (d) 8 million iterations.

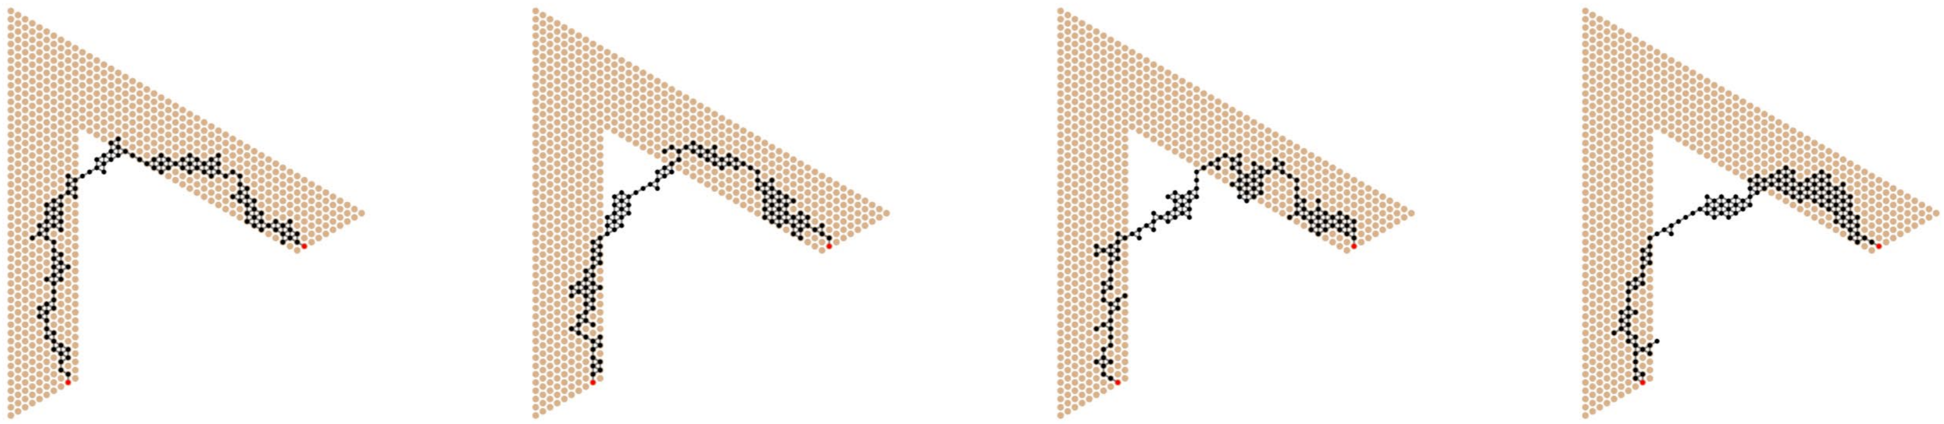

(a) (b) (c) (d)

Shortcut Bridging: $\lambda = 4$, $\gamma = 2$

A particle system initially fully on an N-shaped land mass after 10 million and 20 million iterations.

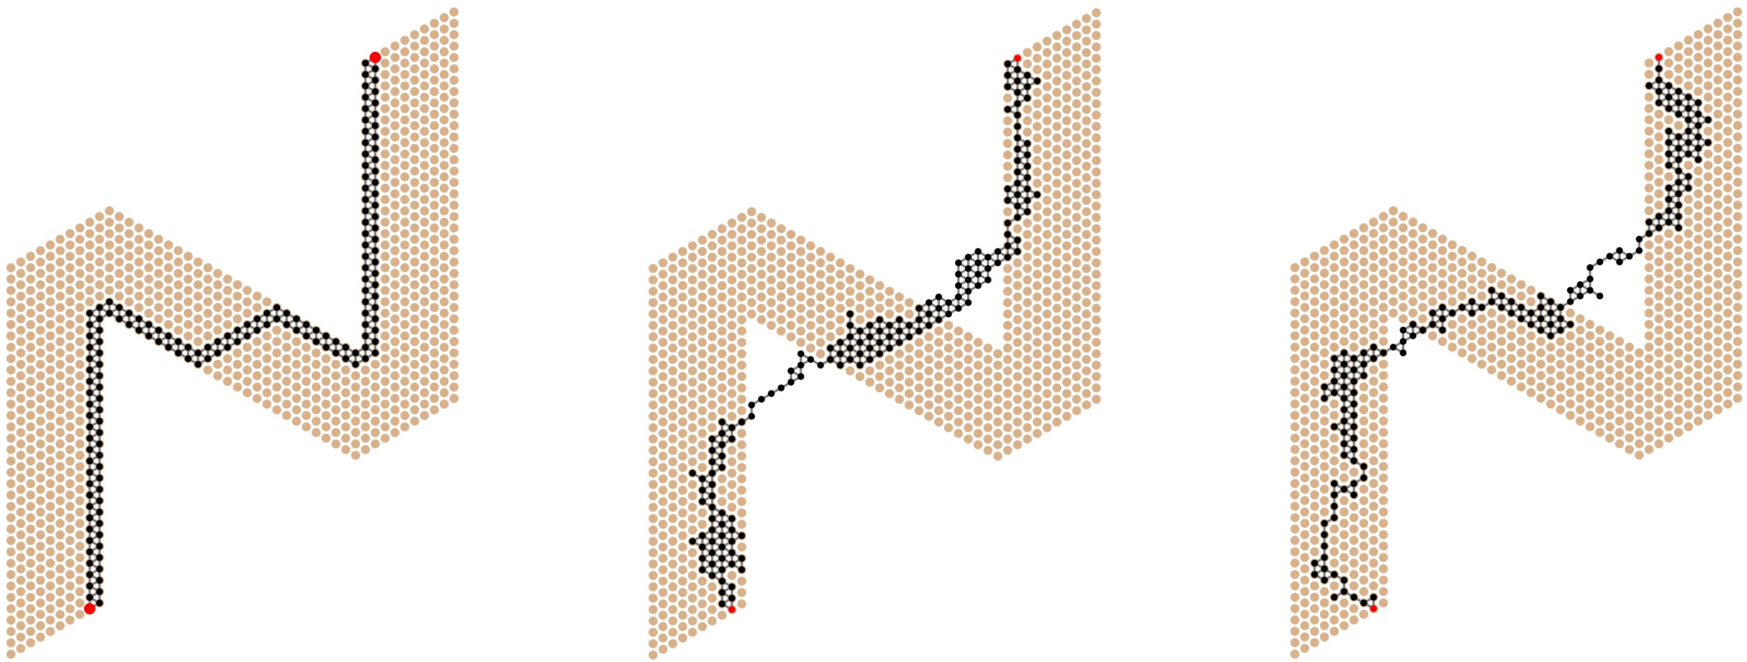

Dependence on Gap Angle

The real army ants' bridges form different shapes depending on the **gap angle**, according to the tradeoff between **shorter paths** and using **too many bridge ants**.

Our algorithm provably exhibits similar behavior: the bridge forms furthest from land on small-angled land mass and hardly even leaves land on the large-angled land mass.

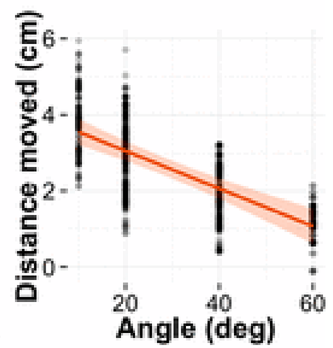

RLPKCG 2015: "Army ants dynamically…"

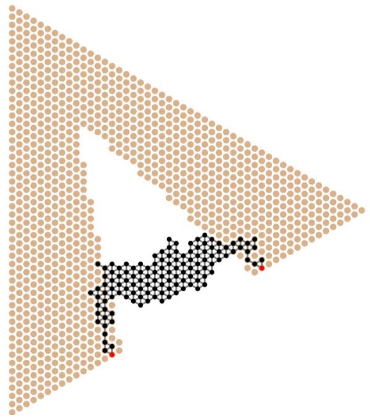

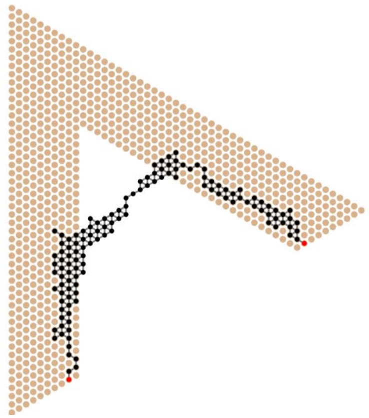

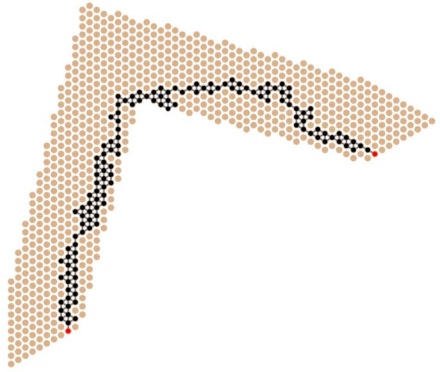

Separation

Problem Statement: Enable a heterogeneous particle system to dynamically separate into large monochromatic clusters or integrate, becoming well-mixed.

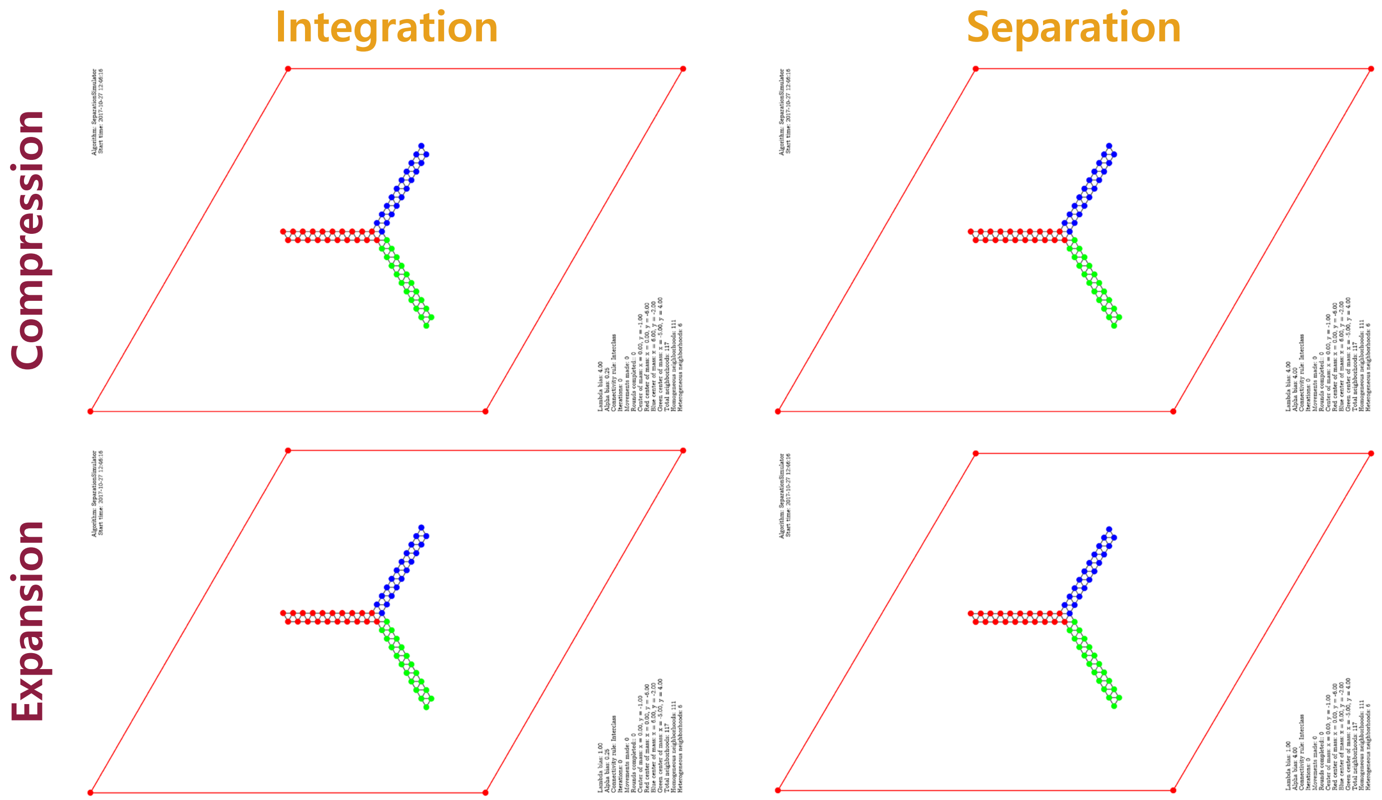

Advantages of the Stochastic Approach

- The algorithms are completely **decentralized** (no leader is necessary for coordination).
- The algorithms **self-stabilize** in the presence of particle **failures**.
- The algorithms are **nearly oblivious**: each particle only keeps 1 bit of memory.

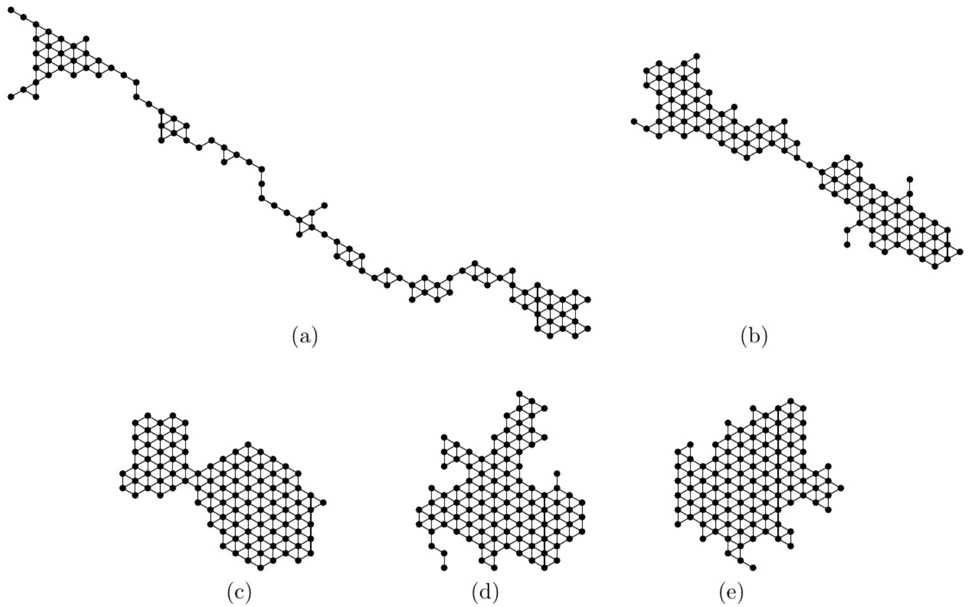

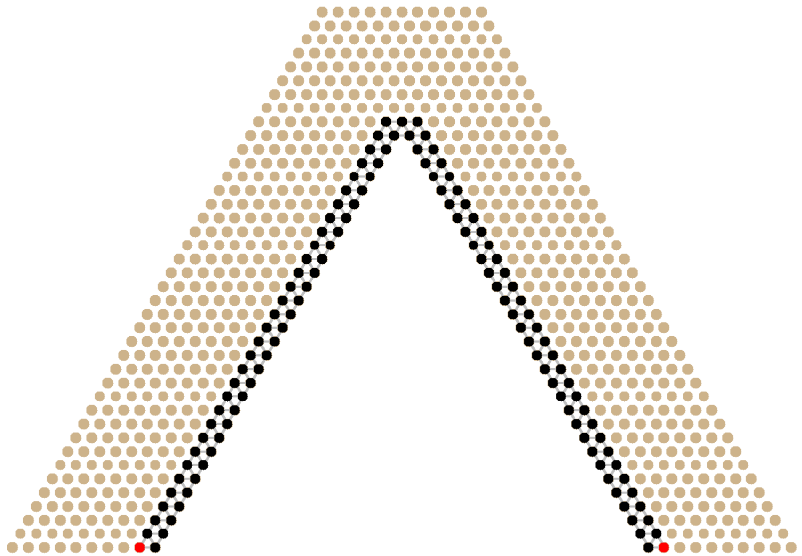

Takeaways of the Stochastic Approach

Good candidate problems for the stochastic approach to programmable matter should:

- Express desired behavior as optimizing a global energy function. For example, in shortcut bridging:

  minimize ***total perimeter*** *and minimize* ***gap perimeter*** -> $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)} \boldsymbol{\gamma}^{-g(\sigma)}$.

- Be able to compute changes in the global energy function using only local information. For example, in compression:

  $\pi(\sigma) \sim \boldsymbol{\lambda}^{-p(\sigma)}$ -> move with probability $\min\{1, \boldsymbol{\lambda}^{\Delta e}\}$.

Our Team

Arizona State University

Georgia Tech

Berkeley UNIVERSITY OF CALIFORNIA

UNIVERSIDAD DE GRANADA

Andréa W. Richa

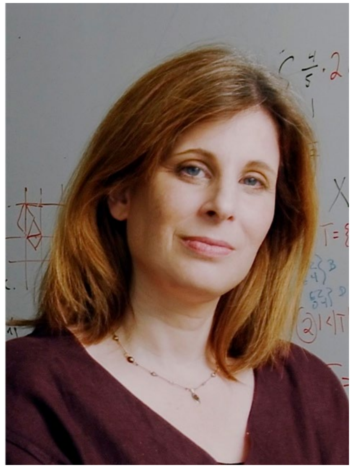

Dana Randall

Cem Gökmen

Sarah Cannon

Marta Andrés Arroyo

Thank you!

sops.engineering.asu.edu

joshdaymude.wordpress.com

Arizona State University