Brief Announcement: Foraging in Particle Systems via Self-Induced Phase Changes

Shunhao Oh ✉
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** ——

The *foraging problem* asks how a collective of particles with limited computational, communication and movement capabilities can autonomously compress around a food source and disperse when the food is depleted or shifted, which may occur at arbitrary times. We would like the particles to iteratively self-organize, using only local interactions, to correctly *gather* whenever a food particle remains in a position long enough and *search* if no food particle has existed recently. Unlike previous approaches, these search and gather phases should be self-induced so as to be indefinitely repeatable as the food evolves, with microscopic changes to the food triggering macroscopic, system-wide phase transitions. We present a stochastic foraging algorithm based on a phase change in the fixed magnetization Ising model from statistical physics: Our algorithm is the first to leverage *self-induced phase changes* as an algorithmic tool. A key component of our algorithm is a careful token passing mechanism ensuring a dispersion broadcast wave will always outpace a compression wave.

1 The Foraging Problem

Collective behavior of interacting agents is a fundamental, nearly ubiquitous phenomenon across fields, reliably producing rich and complex coordination. Examples at the micro- and nano-scales include coordinating cells (including our own immune system or self-repairing tissue and bacterial colonies), micro-scale swarm robotics, and interacting particle systems in physics; at the macro scale it can represent flocks of birds, coordination of drones, and societal dynamics such as segregation. Common properties of many of these disparate systems is that they 1) respond to simple environmental conditions and 2) undergo *phase changes* as parameters of the systems are slowly modified, allowing collectives to gracefully toggle between two often dramatically different macroscopic states.

In the *foraging problem*, we consider a collective of "ants" (i.e., *particles*) with limited computational, communication and movement capabilities that reside on the triangular lattice, along with a *food* particle (i.e., any resource in the environment, e.g., an energy source) that may be placed at any point, removed, or shifted at arbitrary times, possibly adversarially. We would like the particles to consistently self-organize, using only local

interactions, such that if a food particle remains in a position long enough, the particles should transition to a *gather phase* in which many collectively form a single large component with small perimeter around the food. Alternatively, if no food particle has existed recently, the particles should undergo a self-induced phase change and switch to a *search phase* in which they distribute themselves randomly throughout the lattice region to search for food. Unlike previous approaches, this process should be indefinitely repeatable, withstanding overlapping waves of phase changes that may interfere with each other. Like a physical phase change, microscopic changes such as the deletion or addition of a single food particle should trigger these macroscopic, system-wide transitions. This foraging problem has several fundamental application domains, including search-and-rescue operations in swarms of nano- or micro-robots; health applications (e.g., a collective of nano-sensors that could search for, identify, and gather around a foreign body to isolate or consume it, then resume searching, etc.); and finding and consuming/deactivating hazards in a nuclear reactor or a minefield.

2   Model and Preliminaries

In this work, we consider an abstraction of a self-organizing particle system (SOPS), where particles sit on vertices of a finite region of the triangular lattice. We assume particles have constant-size memory, but lack global orientation or any other global information beyond a common chirality. Particles communicate by sending tokens to their nearest neighbors in the lattice, where a *token* is a constant-size piece of information. Individual particles are activated according to their own Poisson clocks, possibly with different rates, and perform instantaneous actions upon activation. Particles are aware of their own and their neighbors' current states and when a particle is activated, it may do a bounded amount of computation, send at most one token (not necessarily identical) to each of its neighbors, and choose one of its six neighbors in the lattice to see if it is unoccupied and move there.[1]

Cannon et al. [1] introduced a related non-adaptive compression and expansion algorithm based on an input parameter λ that defines *system-wide behavior*. When λ is sufficiently small, the system is in an expansion phase, desirable to search for food, while when λ is large, the system will be in a compression phase, desirable when food has been discovered.[2] More specifically, using insight from the Ising model in statistical physics, the authors proved that adding a ferromagnetic attraction λ between particles suffices to stochastically lead the particles in a SOPS to an α-*compressed* configuration with high probability, where the constant $\alpha > 1$ determines an upper bound on the ratio of the configuration perimeter by the minimum possible system perimeter. The Markov chain is defined so that each configuration σ appears with probability $\pi(\sigma) = \lambda^{|E(\sigma)|}/Z$ at stationarity, where $|E(\sigma)|$ is the number of edges in σ and Z is the normalizing constant. It is rigorously shown that the SOPS will reach an α-compressed configuration at stationarity, for some constant $\alpha > 1$, if the attraction force λ is strong enough. Moreover, it is also shown that when the attraction forces are small, the configurations will nearly maximize their perimeter and *disperse* if particles are allowed to disconnect [2], as we do in our algorithm.

Our challenge here is to *self-induce* these system-wide behaviors upon the discovery or depletion of a single food particle. When food is not present, particles communicate to transition to the search phase by collectively lowering λ, and when food is discovered they transition to the gather phase, collectively raising λ to compress around the food.

---

1   Our model can be seen as an abstraction of the (canonical) *Amoebot model* under a *sequential scheduler*.
2   A similar algorithm for the more general setting where particles are allowed to disconnect also provably exhibits a bifurcation, but the notion of compression becomes more complicated [2].

3   The Adaptive Foraging Algorithm

We present the first rigorous local distributed algorithm for solving the foraging problem, the *Adaptive α-Compression algorithm.* There are two main (micro-level) states each particle can be in at any point in time, *dispersion* or *compression*, corresponding to the macro-level *search* and *gather* phases respectively. To switch to the search phase, particles are induced to collectively transition to the dispersion state. Likewise, to switch to gather, particles are induced to transition towards compression. Particles in the dispersion state move around in a process akin to a simple exclusion process, where they perform a random walk while avoiding two particles occupying the same site. Particles enter a compression state when food is found and this information is propagated in the system, resulting in the system gathering and forming a low-perimeter cluster (compressing) around the food. We prove the following:

▶ **Theorem 1.** *Starting from any valid configuration, in the presence of a single food particle that remains static for a sufficient amount of time, the Adaptive α-Compression algorithm will converge to an α-compressed configuration, for any $\alpha > 1$, connected to the food particle at stationarity with high probability. Conversely, if there are no food particles in the system for a sufficient amount of time, the system converges to a uniform distribution over all possible assignments of particles to sites on the lattice.*

We believe Adaptive α-Compression is the first adaptive algorithm to leverage a *self-induced phase change as an algorithmic tool.* The challenge is to share information locally and autonomously so that eventually most particles enter the correct state and the system exhibits the appropriate phase behavior. We rely on token passing for the system to be able to collectively transition between (multiple, possibly overlapping and interfering) gather and search phases: Each particle locally adjusts its ferromagnetic bias parameter λ to be high when it receives compression tokens, which are continuously generated by any particle in contact with the food source, and to be low when it receives dispersion tokens, which are flooded through the network once a food particle disappears. In order to ensure that, our token passing scheme needs to be carefully engineered so that when the food particle moves or vanishes, the *rate at which the compressed cluster around the food dissipates (via particles returning to the dispersion state) outpaces the rate at which the cluster may continue to grow (via particles joining the cluster in a compression state)*, and thus that the broadcast wave of dispersion tokens will always outpace the broadcast wave of compression tokens, ensuring that whenever we have a situation where two phase change waves compete, the dispersion wave will be the one which wins out in the end. This is done via a novel potential function argument that carefully sets the dispersion and compression token passing probabilities.

We note that while Adaptive α-Compression is very similar to the non-adaptive compression algorithm [1] in the presence of food, allowing particles to compress around a single fixed point (the food particle), this is a nontrivial generalization. Even proving ergodicity of the underlying Markov Chain in the presence of a fixed (food) point from which other particles cannot disconnect is quite complicated and does not follow directly from [1].

───── **References** ─────────────────────────────────

1   Sarah Cannon, Joshua J. Daymude, Dana Randall, and Andréa W. Richa. A Markov chain algorithm for compression in self-organizing particle systems. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, pages 279–288, 2016.

2   Shengkai Li, Bahnisikha Dutta, Sarah Cannon, Joshua J. Daymude, Ram Avinery, Enes Aydin, Andréa W. Richa, Daniel I. Goldman, and Dana Randall. Programming active granular matter with mechanically induced phase changes. *Science Advances*, 7(17):eabe8494, 2021.