Low-Communication Distributed Optimization via E. Coli Swarm Foraging

Shashank Singh [1]   Saket Navlakha [2]   Ziv Bar-Joseph[1]

2nd Workshop on Biological Distributed Algorithms

October 12, 2014

[1]Carnegie Mellon University, Pittsburgh, PA, USA
[2]Salk Institute for Biological Studies, La Jolla, CA, USA

Differences from Insect Foraging

Insect Colonies	Bacteria Swarms
agents move food to colony	swarm moves to food
fixed pheromone trails	diffusing protein signals
nurses, foragers, queen, etc.	identical cells
complex navigation abilities	no navigation ability

Bacteria Swarm Foraging

- Food source which diffuses with density $f : \mathbb{R}^2 \to \mathbb{R}$ throughout solution
- Obstacles
- Bacteria swarms (typically 1-4 swarms of 20-50 agents each)

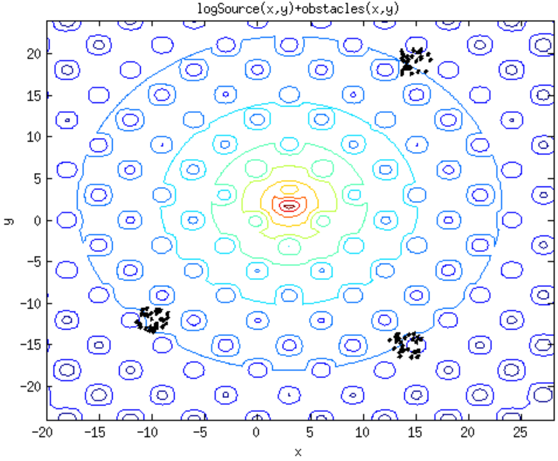

Computational Problem

Several nodes each want to maximize the same objective function:

$$\max_{x \in S \subseteq \mathbb{R}^d} f(x).$$

- can evaluate f, but don't know its form

Computational Problem

Several nodes each want to maximize the same objective function:

$$\max_{x \in S \subseteq \mathbb{R}^d} f(x).$$

- can evaluate f, but don't know its form
- S and f typically non-convex
  - Can have small local maxima

Computational Problem

Several nodes each want to maximize the same objective function:

$$\max_{x \in S \subseteq \mathbb{R}^d} f(x).$$

- can evaluate f, but don't know its form
- S and f typically non-convex
  - Can have small local maxima
- Individual nodes computationally weak

Computational Problem

Several nodes each want to maximize the same objective function:

$$\max_{x \in S \subseteq \mathbb{R}^d} f(x).$$

- can evaluate f, but don't know its form
- S and f typically non-convex
  - Can have small local maxima
- Individual nodes computationally weak
- Nodes can broadcast (small) messages to nearby nodes

Individual Movement (Tumbling)

Each iteration, each agent perturbs its direction based on previous change in food density:

$$\delta = f(x_t, y_t) - f(x_{t-1}, y_{t-1})$$

$$\sigma \propto \max\{0, 1-\delta\}.$$

$$\theta \to \theta + \varepsilon, \quad \text{where} \quad \varepsilon \sim \mathcal{N}(0, \sigma^2),$$

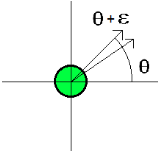

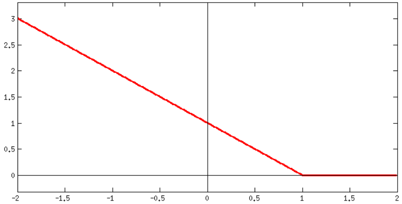

Individual Movement (Tumbling)

This works, but very inefficiently:

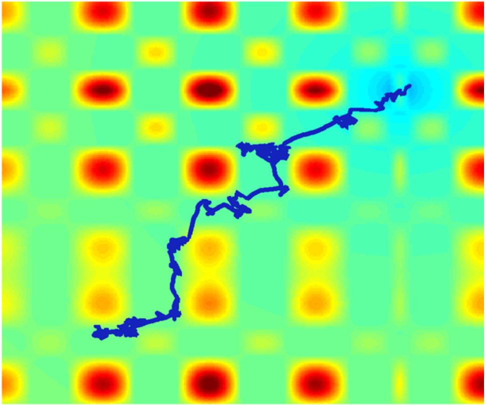

Basic Swarm Movement (Shklarsh et al., 2011)

On each iteration, each agent combines its (perturbed) velocity with the influence of the swarm

$$v_{i,t+1} = w_v R_\varepsilon v_{i,t} + \begin{cases} w_r r_{i,t} & \text{if any neighbors are too close} \\ w_a a_{i,t} + w_\omega \omega_{i,t} & \text{else} \end{cases}$$

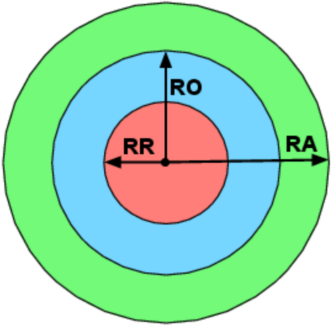

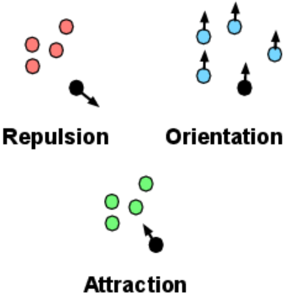

Basic Swarm Movement (Repulsion)

Avoid collisions and spread out to cover area

$$r_{i,t} = \sum_{x_{j,t} \in B_{RR}(x_i)} \frac{x_{j,t} - x_{i,t}}{\|x_{j,t} - x_{i,t}\|}.$$

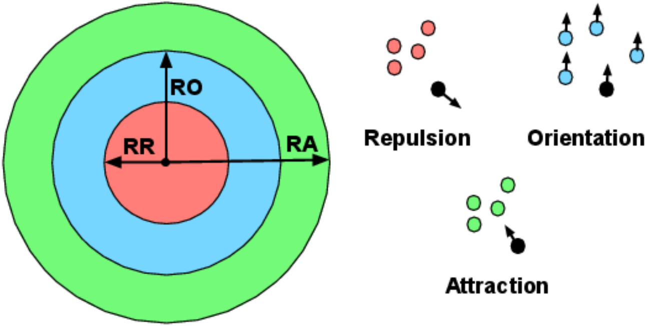

Basic Swarm Movement (Attraction)

Stay together as a group

$$a_{i,t} = \sum_{x_{j,t} \in B_{RA}(x_i) \setminus B_{RO(x_i)}} \frac{x_{j,t} - x_{i,t}}{\|x_{j,t} - x_{i,t}\|}.$$

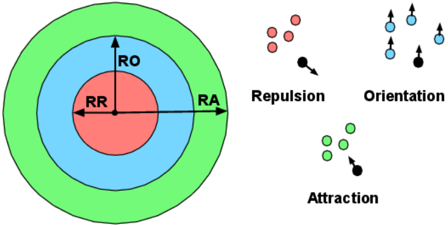

Basic Swarm Movement (Orientation)

Move similarly to your neighbors

$$\omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}.$$

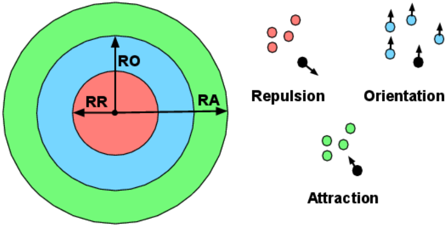

Basic Swarm Movement (Orientation)

Move similarly to your neighbors

$$\omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}.$$

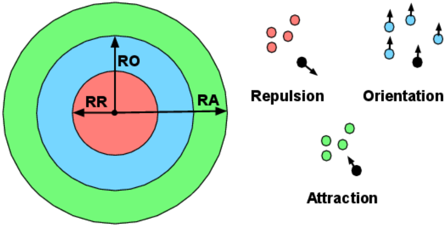

- Accelerates swarm when the correct direction is clear

Basic Swarm Movement (Orientation)

Move similarly to your neighbors

$$\omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}.$$

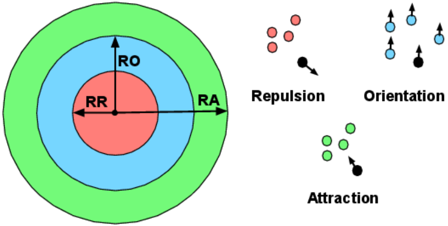

- Accelerates swarm when the correct direction is clear
- Helps "smooth" interactions by preventing collisions.

Basic Swarm Movement (Shklarsh et al.)

Again,

$$v_{i,t+1} = w_v R_\varepsilon v_{i,t} + \begin{cases} w_r r_{i,t} & \text{if any neighbors are too close} \\ w_a a_{i,t} + w_\omega \omega_{i,t} & \text{else} \end{cases}$$

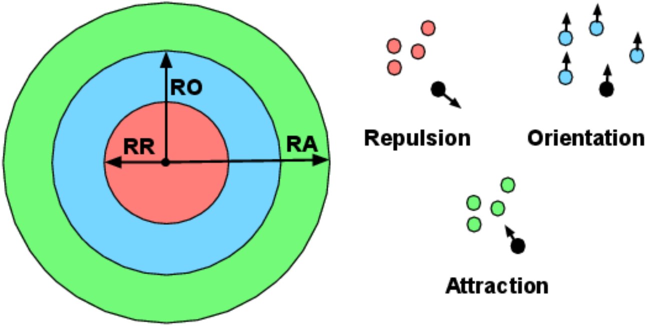

Issues

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction.

Issues

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction.

$$a_{i,t} = \sum_{x_{j,t} \in B_{RA}(x_i) \setminus B_{RO(x_i)}} \frac{x_{j,t} - x_{i,t}}{\|x_{j,t} - x_{i,t}\|} \quad \text{and} \quad \omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}$$

- Messages can be continuous (e.g., floats)

Issues

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction.

$$a_{i,t} = \sum_{x_{j,t} \in B_{RA}(x_i) \setminus B_{RO(x_i)}} \frac{x_{j,t} - x_{i,t}}{\|x_{j,t} - x_{i,t}\|} \quad \text{and} \quad \omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}$$

- Messages can be continuous (e.g., floats)
    - Real bacteria send protein signals of only a few bits

Issues

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction.

$$a_{i,t} = \sum_{x_{j,t} \in B_{RA}(x_i) \setminus B_{RO(x_i)}} \frac{x_{j,t} - x_{i,t}}{\|x_{j,t} - x_{i,t}\|} \quad \text{and} \quad \omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}$$

- Messages can be continuous (e.g., floats)
  - Real bacteria send protein signals of only a few bits
- Receiver's measurements can be arbitrarily large

Issues

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction.

$$a_{i,t} = \sum_{x_{j,t} \in B_{RA}(x_i) \setminus B_{RO(x_i)}} \frac{x_{j,t} - x_{i,t}}{\|x_{j,t} - x_{i,t}\|} \quad \text{and} \quad \omega_{i,t} = \sum_{x_{j,t} \in B_{RO}(x_i)} \frac{v_{j,t}}{\|v_{j,t}\|}$$

- Messages can be continuous (e.g., floats)
  - Real bacteria send protein signals of only a few bits
- Receiver's measurements can be arbitrarily large
  - Real bacteria distinguish only a few levels

Discretization and Thresholding

- Introduce a thresholding discretization function:
  - For $T > 0$, $L \in \mathbb{N}$, $\|D_{L,T}(x)\| = \min\{T, \lfloor L\|x\|\rfloor / L\}$.
  - Approximate vectors by cardinal vectors to discretize direction

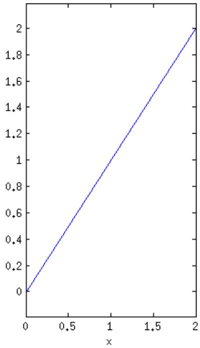

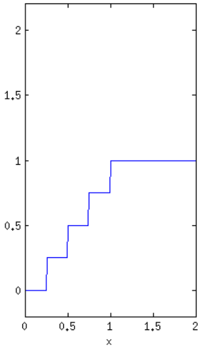

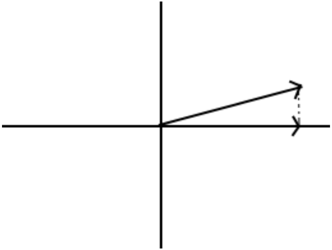

Issues (Cont.)

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction (repulsion is ok).

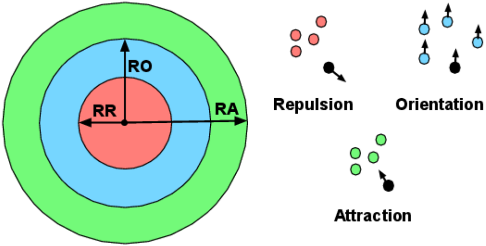

- Agents can identify message senders (dedicated channels)
  - Requires $\log(n)$ extra bits per message
  - Swarm can be dynamic
  - Real bacteria broadcast to their neighbors

Issues (Cont.)

The Basic Swarm Movement model makes unrealistic assumptions about how bacteria communicate orientation and attraction (repulsion is ok).

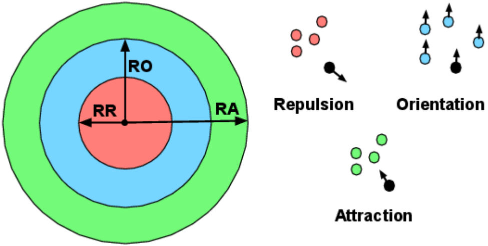

- Agents can identify message senders (dedicated channels)
  - Requires $\log(n)$ extra bits per message
  - Swarm can be dynamic
  - Real bacteria broadcast to their neighbors
- Ability to communicate is unaffected by distance

Distance Weighting

- Broadcast messages, but weight communication by distance
  - Messages decay exponentially with distance:
    $w_a(x) = \exp(-c_a x), \quad w_\omega(x) = \exp(-c_\omega x) \quad (c_\omega > c_a)$

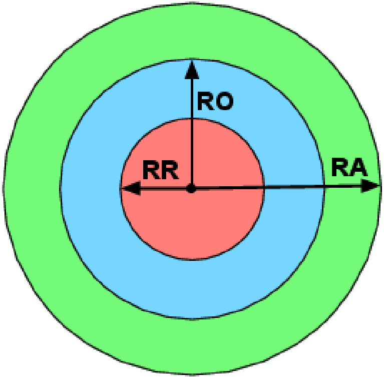

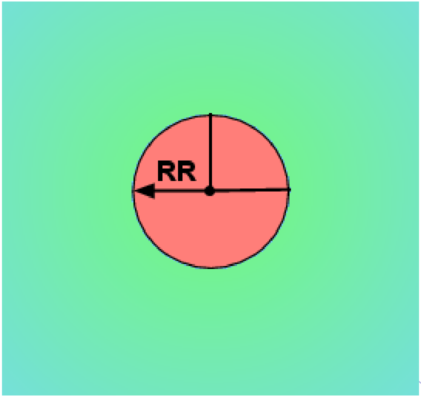

Efficient Communication Model

- Discretize after weighting:

$$a_{i,t} = \sum_{j=1}^{n} D_{L,T}\left(w_a(\|x_j - x_i\|)\frac{(x_j - x_i)}{\|x_j - x_i\|}\right)$$

$$\omega_{i,t} = \sum_{j=1}^{n} D_{L,T}\left(w_a(\|v_{j,t}\|)\frac{v_j}{\|v_j\|}\right)$$

Recall

$$v_{i,t+1} = w_v v_{i,t} + \begin{cases} w_r r_{i,t} & \text{if any neighbors are too close} \\ w_a a_{i,t} + w_\omega \omega_{i,t} & \text{else} \end{cases}$$

Experimental Results

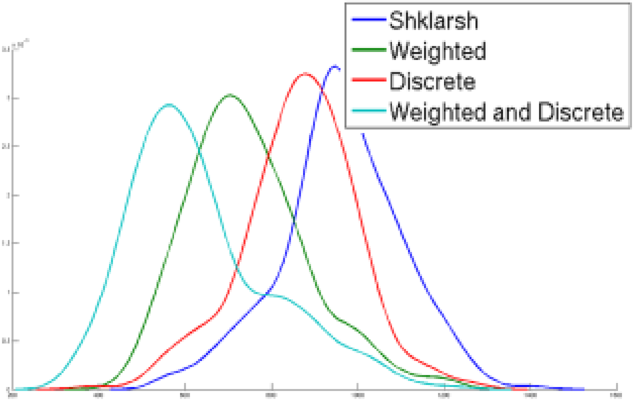

Path Length

Adaptive Listening

Help if you're making progress, get help if you're stuck

- weight current velocity based on performance

Modified model:

$$v_t = w(\delta) \cdot v_{t-1} + (1 - w(\delta))u,$$

where w is increases with $\delta = f(x_t, y_t) - f(x_{t-1}, y_{t-1})$.

Silent Agents

- broadcasting messages takes energy
- many messages are redundant
- under scarce resources, may not want to help competition

Silent Agents

- broadcasting messages takes energy
- many messages are redundant
- under scarce resources, may not want to help competition

Modified model: For some $p_s \in [0, 1]$, each agent is silent with probability p_s.

Experimental Results: Silent Agents

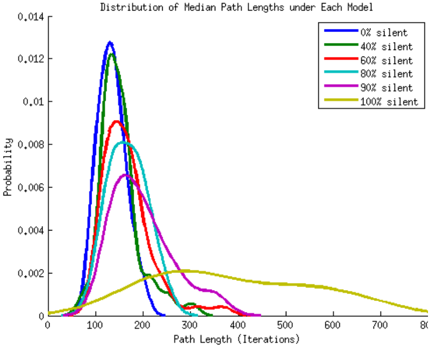

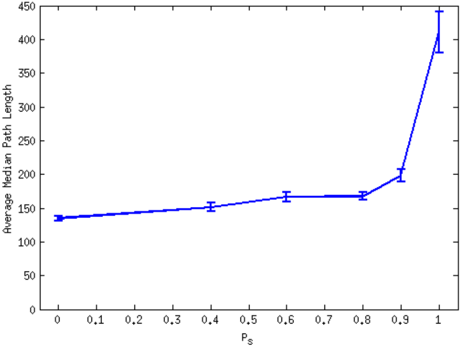

Very few agents actually need to communicate!

Summary

- Primitive bacteria solve computationally challenging problems collectively
- Swarm communication is helpful even under highly restricted communication
  - Agents need only broadcast a few bits
  - Signals only need need to travel short distances
  - Only some agents need to communicate

Future Work

- Consider competition (finite food sources)
- Multiple food sources/mixed objectives
  - Agents can have different preferences
- Compare to biological model
  - Can identify genes responsible for communication?
  - How is orientation really communicated?
- Theory
  - Convergence rates
  - Lower bounds

Thanks!

Simulation code is available on GitHub.