

Behavior Modes for Randomized Robotic Coverage

Jacob Beal
BBN Technologies
Cambridge, MA 02138
Email: jakebeal@bbn.com

Nikolaus Correll, Leonardo Urbina
MIT CSAIL
Cambridge, MA 02139
Email: nikolaus@csail.mit.edu, lurbina@mit.edu

Jonathan Bachrach
Makani Power
Alameda, CA 94501

Abstract—A basic primitive in a networked robotic swarm is to form a connected component that covers some area with relatively uniform density. Although most approaches to the problem require local coordinate information, it has been proposed that robots with only connectivity information do this instead with a generalized form of diffusion-limited aggregation, in which robots wander randomly until they find a location where their topological constraints are satisfied and they are connected to a designated seed point. We find that the behavior of the algorithm varies qualitatively along a spectrum defined by the relative size of the total area, covered area, and initial distribution of robots. We identify and analyze five representative behaviors along this spectrum, finding that fast convergence can be expected only within a small range of parameters. Further, our results suggest that general coverage algorithms may require that the robotic swarm be coordinated across long distances.

I. MOTIVATION

A basic primitive in a networked robotic swarm is to form a connected component that covers some area with relatively uniform density. Applications that demand this task include exploring a confined space for search and rescue (e.g. [1]), deploying wireless coverage (e.g. [2]), and scouting for urban combat (e.g. [3]). Once distributed, the swarm may be largely stationary, as in the cases above, or may attempt to move as an aggregate. In this paper we will consider only the stationary case, which we believe is informative for the mobile case as well.

Many approaches to this problem require local coordinate information, such as distance and angle to neighbors or relative heading (see, for example, [4] and [5]). Including sensors to acquire such information may be infeasible, however, particularly when considering the construction of a swarm containing a large numbers of tiny robots.

One algorithm[6] to form a stationary connected component without such information is a generalized form of diffusion-limited aggregation[7]. Robots wander randomly until they find a location where their number of neighbors is within an acceptable band, neither too low nor too high,¹ and they are connected through the network to a designated seed point. This algorithm is attractive because it is simple, requires little communication, and does not require any sort of localization on the part of the robots. But can random wandering produce good enough behavior?

¹In pure DLA, there is no “too high” case and robots never restart once they have stopped.




Fig. 1. Robots are satisfied and stationary (green) with a moderate number of neighbors and connection to the seed (purple). Robots are unsatisfied (red) if the number of neighbors is too high (dark) or too low (medium), or there is no path to the seed (light). The grey background shows the union of the robots’ unit disc communication areas. This illustration shows robots with $\alpha = 3$ and $\beta = 5$.

A previous analysis of DLA by Poduri and Sukhatme indicates that the swarm should converge in $O(1/\sqrt{n})$ time for n robots, but only when assuming that the location covered by a robot in each random step is independent of its previous history and that the distribution of robots through space remains uniform even once some begin to connect and stop moving[8], [9]. These assumptions, however, approximate behavior well only in the special case where the final area covered by the robots is close to that of their initial distribution and to the total area of the environment. We show that the general convergence time is much slower.

We find instead that the behavior of the more general algorithm (and its DLA special case) varies qualitatively along a spectrum defined by the relative size of the total area, covered area, and initial distribution of robots. We identify and analyze five representative behaviors along this spectrum, finding that fast convergence can be expected only within a small range of parameters. Further, our results suggest that fast coverage may always require coordination across long distances.

II. MODEL

We use the following model for the robots and their environment. Some assumptions are not used by the algorithm or our analysis, but serve to illuminate the design space in which the algorithm is situated.

- Robots communicate wirelessly using a unit disc model of radius r : every pair of robots within r meters of one

```

For robot  $i$ :
IF  $\alpha \leq |nbrs(i)| \leq \beta$ 
AND connected to seed
THEN do nothing
ELSE move  $v \cdot W_t$ 
      (where  $W_t$  is a Wiener process)

```

Fig. 2. Pseudocode for randomized coalescence algorithm.

another are neighbors in the network graph.

- Robots can move freely through a static, connected area of two-dimensional space A .
- One robot (or non-robotic node) is designated as the seed, and does not move. All n others are initially distributed through an area $D \subseteq A$, using a uniform random distribution.
- Communication is much faster than robot motion, enough so that for purposes of analysis, we will assume communication is instantaneous and perfect.
- Robots have no information about how many other robots there are or where they are located (either locally or globally), only the identities of their neighbors. Robots can uniquely distinguish one another.
- Mechanical constraints of robot movement (e.g. turning speed) are assumed to not be limiting factors.

III. RANDOMIZED COALESCENCE ALGORITHM

The randomized coalescence algorithm we use, from [6] is shown in Figure 2. Letting α be the minimum acceptable number of neighbors and β be the maximum acceptable number of neighbors, a robot stops moving if it has an acceptable number of neighbors and is part of a connected component that includes the seed. Otherwise, the robot moves using a Wiener process multiplied by v , a velocity constant. A Wiener process is a continuous-time stochastic process where the difference in values across any elapsed time period t is Gaussian with mean 0 and variance t , and non-overlapping time periods are independent. This provides scale-invariant random wandering, such as Brownian motion, and may be approximated by a random walk. Figure 1 shows the various cases for the algorithm, with some robots satisfied and stationary and others unsatisfied and mobile.

Diffusion-limited aggregation may be viewed as a special case of this algorithm, in which $\alpha = 1$, $\beta = \infty$. The coalescence studies in [8] and [9] are a special case as well, although they introduce the possibility of clustering away from the seed. Note that under the more general algorithm we use, a stopped robot may start moving again if it loses or gains neighbors, or if motion elsewhere disconnects it from the seed; robots never begin moving again in the two previously studied special cases.

For purposes of this paper, we assume that $\beta > 2$ and that $\beta > \alpha$. If $\beta \leq 2$ it is not possible for the aggregate to branch, meaning the algorithm can only converge if the robots form a line, and if $\beta = \alpha$, then incremental convergence is not possible.

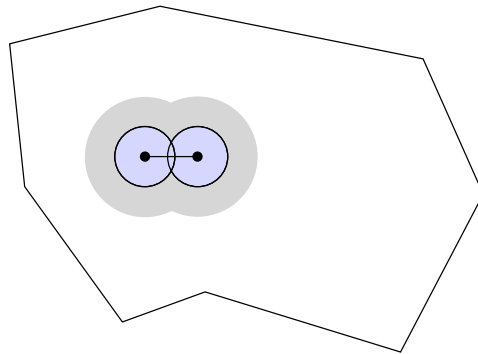


Fig. 4. If we surround each robot (black dots) with a disc half the radius of the robot’s communication coverage, then if two discs overlap at any point, it implies that the robots can communicate.

IV. BEHAVIORAL SPECTRUM

Although there are seven variables all told, the behavior of the system is controlled by the relationship between three areas: the area A through which robots can move, the area D through which robots are initially distributed, and the range of areas $[T, L]$ from the area T covered by the tightest possible converged arrangement of robots to the area L covered by the loosest. All the parameters are encapsulated in this relationship: the coverage areas are derived from n , r , and α for L and β for T , while the velocity constant v scales convergence without changing it qualitatively.

For simplicity, we further assume that A and D are both fairly open shapes where a random walk can easily move a node from one part of the space to another. In more tortuous areas, such as two open spaces connected by a narrow “neck,” convergence may be much slower.

At different relations between A , D , and $[T, L]$, different phenomena dominate, as shown in Figure 3. When the area is too small, finding any arrangement where the upper density limit is satisfied for all robots becomes difficult or impossible. When the robots easily fit within A , the initial distribution becomes important, and either the search for the seed (e.g. robots converging after an air-drop) or diffusion away from the source (e.g. robots fanning out to do search and rescue) may dominate. Only when these two are at relatively equal strengths is it possible for convergence to be rapid. The remainder of the paper presents these five representative cases and our analysis of their convergence rates.

V. CONVERGENCE AT HIGH DENSITY

If there is any robot with more than β neighbors, the algorithm has not converged, since that robot will continue wandering randomly. If we consider discs of $1/2$ the radius r that a robot’s communication covers, then whenever two discs overlap, it implies that the robots at the source of the discs are neighbors (Figure 4). Thus, if we have a point which is k -covered by half-radius discs, then it implies that every robot covering that point has at least $k - 1$ neighbors.

If we neglect edge effects (which are small for large areas), then each robot’s half-disc covers a region of size $\frac{1}{4}\pi r^2$. If

| Name | Condition | Convergence | Dominant Phenomena |
|------------------|--------------------------|--------------------|--|
| Impossible | $A < T$ | Impossible | Density too high |
| Critical | $A \approx T$ | Arbitrarily long | Unlikely perfect arrangement |
| Discovery | $A \gg D \gg L$ | $\Omega(D)$ | Difficulty in finding seed. |
| Fast Aggregation | $A \gg D \approx [T, L]$ | $\Theta(\sqrt{n})$ | “Crystallization” from seed (unproved) |
| Dispersion | $A \gg T \gg D$ | $\Omega(n)$ | Diffusion from initial distribution |

Fig. 3. Convergence times for five representative conditions along the behavioral spectrum, and the dominant phenomena that causes that convergence time. A is the total area, D is the area robots are initially distributed through, and T and L are the areas of the tightest and loosest possible areas covered when the algorithm has converged.

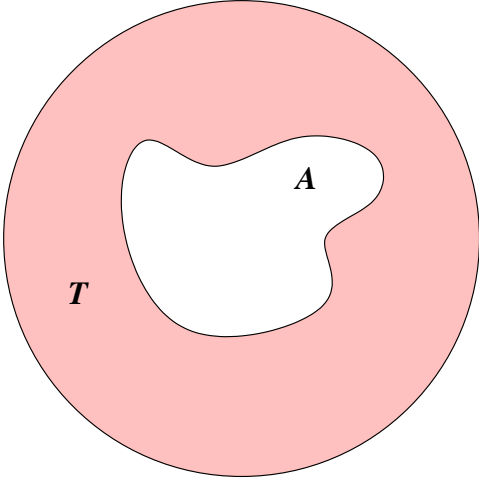


Fig. 5. In the *impossible* case, the area A through which the robots can move is smaller than the area T covered by robots in their tightest allowable distribution. In this case, the robots will never converge.

the total area covered goes above $A(\beta + 1)$, then some point must be covered more than $\beta + 1$ times, meaning that some robot has more than β neighbors. A group of n robots covers an area of $\frac{1}{4}\pi r^2 n$, so we have a bound, above which there can be no convergence, of

$$\frac{1}{4}\pi r^2 n \leq A(\beta + 1)$$

$$\frac{\pi r^2 n}{4(\beta + 1)} \leq A$$

This is also a lower bound on the area T covered by the tightest possible converged arrangement of robots. We name this the *impossible* case (Figure 5).

Edge effects can be accounted for by adding points near the perimeter that can be covered. The amount of space to be added depends on whether the boundary affects robots ability to communicate with one another (Figure 6). If it blocks communication, as might a wall, then the area is effectively larger than if it permits communication, as might a boundary caused by a pit. In no case, however, is the effective area increased by more than $\frac{1}{4}\pi r^2 P$, where P is the perimeter of the area. For most spaces of interest in this problem, this is likely to be a small effect, as the area of the interior will be much larger than the area near the perimeter.

Note that this is still not a tight bound, because it neglects the geometric constraints of sphere packing and the shape

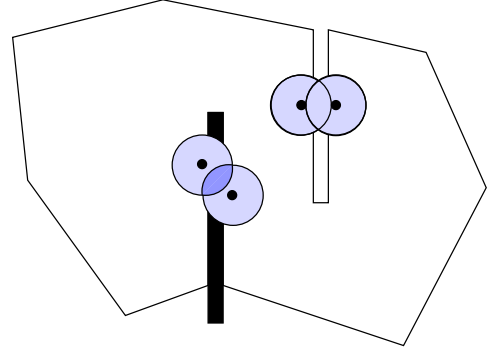


Fig. 6. A tighter bound on packing can be found by adding in area near the perimeter of the space. The amount of area to be added depends on whether the boundary of the space blocks communication between robots (left pair), as might a wall, or permits communication (right pair), as might a boundary caused by a pit.

of the area. That effect will change things by only a small constant multiple, however.

If the area A is increased slightly, such that convergence is possible, the time to convergence can be made arbitrarily high. An extremely loose lower bound on convergence time can be found by considering the ratio between the total area and the area of points where a robot can stop when the algorithm has converged. By setting the area arbitrarily close to the threshold of impossibility, the “wiggle room” for robots can be reduced arbitrarily close to zero, guaranteeing arbitrarily long convergence time in this *critical* case (Figure 7).

VI. CONVERGENCE AT LOW DENSITY

We now turn instead to the case where the area A in which the robots are wandering is much bigger than the area they can cover. In practice, one would want to weaken the definition of convergence so that the system can be considered to have converged even if a few robots get “lost” far away from the seed, but for the analysis we are going to perform here it will not matter.

First, let us establish an upper bound for the area L covered by the loosest possible distribution. For even α , the distribution of robots that covers the greatest area is a line (neglecting the ends) with $\frac{2r}{\alpha}$ meters between robots (Figure 9)—by the triangle inequality, we can see that for any arrangement of robots, the amount of area covered by the addition of a new robot is maximized by placing the new robot as far as possible from previous robots. For odd α , the optimal distribution is the




Fig. 7. In the *critical* case, the area A through which the robots can move is just barely larger than the area T covered by robots in their tightest allowable distribution. In this case, the time to convergence may be unboundedly high.

same as for $\alpha + 1$, except that every $\alpha + 1$ st interval is doubled, producing the same mean advancement of $\frac{2r}{\alpha}$ per step. Taking a bounding box around this line, we can thus bound L above by

$$L < 2r \cdot n \frac{2r}{\alpha} = n \frac{4r^2}{\alpha}$$

The behavior of the system now depends on the relationship between the initial distribution area D and the range $[T, L]$ of areas where convergence can occur. For simplicity, we will always assume that the initial distribution is centered on the seed and has relatively low eccentricity.

A. Dispersion-Dominated Convergence

When the initial distribution area D is much smaller than the area T of the tightest converged arrangement of robots, the random *dispersion* of the robots outward dominates (Figure 8(a)).

We may assume without loss of generality that the robots begin far from the boundary of A . Let us also assume that their arrangement when converged will be approximately circular, minimizing the expected distance that robots must travel. Even so, since T grows with n , the majority of the robots will need to travel $\Theta(\sqrt{n})$ meters outward from their starting points. Since the robots wander using a two-dimensional Wiener process, the expected time for them to cover a distance \sqrt{n} is n . Thus the convergence time for this case is $\Omega(n)$.

B. Coalescence-Dominated Convergence

When the initial distribution area D is much larger than the area L of the loosest converged arrangement of robots, on the other hand, the dominant cost comes from the robots trying to find the seed. We can find an extremely conservative lower-bound of the convergence time for this *discovery* case (Figure 8(b)) by considering how long it takes for a robot to go far enough that it could even get near the seed.

We assume again that the boundary of A is not involved. Since we assume that D is relatively circular, with the seed at its center, the distance from the robots to the seed is




Fig. 9. For even α , the distribution of robots that covers the greatest area is a line with $\frac{2r}{\alpha}$ meters between robots (neglecting the ends), as in the fragment shown above for $\alpha = 4$. For odd α , the distribution is the same except as for $\alpha + 1$, except that periodically there is a double-length interval.

bounded. Since D is much larger than L , however, the distance from most robots' initial positions to their eventual converged positions is effectively the same as their distance to the seed. A majority of the robots must therefore travel $\Theta(\sqrt{D})$ meters in order to reach their final position. Since the robots wander using a two-dimensional Wiener process, the expected time for them to cover a distance \sqrt{D} is D . Thus the convergence time for this case is $\Omega(D)$.

C. Critical Convergence

When dispersion and discovery are in balance and D is close to the range $[T, L]$, convergence may be much faster. We have not proved the rate, but will sketch our preliminary analysis.

In this *fast aggregation* case (Figure 8(c)), robots are packed at approximately the same density when stationary as when wandering freely, so the surface of the connected component grows at approximately the same rate that wandering robots are depleted from its vicinity, maintaining the uniform distribution (although some fraction of the robots may “escape” into the larger area A). As a result, at any time there are likely to be a large enough number of robots within an approximately constant distance of the surface of the connected component to keep it growing at a linear rate. Since we are considering 2D aggregation, a linear rate of surface growth would result in a convergence time of $\Theta(\sqrt{n})$.

The approximations Poduri and Sukhatme use in [8] and [9] may also apply, implying a tighter $O(1/\sqrt{n})$ bound. This may require A to be close to D as well, or there may be some even tighter constraint that is required to account for their simulation results.

VII. CONTRIBUTIONS

We have demonstrated that the convergence behavior of random coalescence varies qualitatively along a spectrum defined by the relative size of the total area, covered area, and initial distribution of robots. We have further identified and analyzed five representative behaviors along this spectrum, finding that fast convergence can be expected only within a small range of parameters.

Since the slow ranges correspond to reasonable deployment scenarios, such as dispersion from a source or coalescence following an air-drop, this suggests that random motion is a poor choice for driving robot coverage algorithms in the absence of relative coordinate information. But might some




Fig. 8. In the *dispersion* case (a), the area of the initial distribution D is much smaller than the area T covered by robots in their tightest allowable distribution, which is much smaller than the area A through which the robots can move. In this case, the time to convergence is $\Omega(n)$. In the *discovery* case (b), D is much larger than the area L covered by robots in their loosest allowable distribution, which is much smaller than A . In this case, the time to convergence is $\Omega(D)$. In the *fast aggregation* case (d), D is close to the range of areas covered by robots in allowable distributions. In this case, the time to convergence is much faster, perhaps $\Theta(\sqrt{n})$.

more intelligent strategy based on local information—like the integrate-and-tumble *e. coli* chemotaxis strategy—improve performance?

In fact, it appears that any algorithm based on purely local information is likely to fail similarly: coalescing robots need some sort of long-range information in order to find the seed efficiently, and dispersing robots in the middle of the swarm need to know which way is outward. The environment which the robots are navigating provides no information about where the other robots are located (and thus how to navigate), so even sophisticated strategies using only local information cannot be expected to improve significantly on the random algorithm. One way or another the robots must exploit some sort of coordination across long distances if coverage is to be fast enough to be effective.

ACKNOWLEDGMENT

The authors would like to thank the US and Swiss National Science Foundations for their support.

REFERENCES

- [1] V. Kumar, D. Rus, and S. Singh, "Robot and sensor networks for first responders," *IEEE Pervasive Computing*, vol. 3, no. 4, pp. 24–33, 2004.
- [2] D. IPTO, "Landroids," <http://www.darpa.mil/ipto/programs/ld/ld.asp>, Retrieved Nov. 10, 2008.
- [3] J. McLurkin and J. Smith, "Distributed algorithms for dispersion in indoor environments using a swarm of autonomous mobile robots," in *Distributed Autonomous Robotic Systems Conference*, June 2004.
- [4] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.
- [5] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: Algorithms and theory," *IEEE Trans. on Automatic Control*, vol. 51, no. 3, pp. 401–420, March 2006.
- [6] N. Correll, J. Bachrach, D. Vickery, and D. Rus, "Ad-hoc wireless network coverage with robots that cannot localize," in *To appear in IEEE International Conference on Robotics and Automation (ICRA)*, 2009.
- [7] T. A. Witten and L. M. Sander, "Diffusion-limited aggregation, a kinetic critical phenomenon," *Phys. Rev. Lett.*, vol. 47, no. 19, pp. 1400–1403, Nov 1981.
- [8] S. Poduri and G. S. Sukhatme, "Achieving connectivity through coalescence in mobile robot networks," in *International Conference on Robot Communication and Coordination (ROBOCOMM)*, October 2007.
- [9] —, "Latency analysis of coalescence in robot groups," in *IEEE International Conference on Robotics and Automation (ICRA)*, April 2007.