Molins, P., & Hauert, S. (2019). Trail Formation using Large Swarms of Minimal Robots. In *The 34th ACM/SIGAPP Symposium On Applied Computing* (pp. 946-952). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3297280.3297374>

Peer reviewed version

License (if available):
Other

Link to published version (if available):
[10.1145/3297280.3297374](https://doi.org/10.1145/3297280.3297374)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via ACM at <https://doi.org/10.1145/3297280.3297374> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Trail Formation using Large Swarms of Minimal Robots

Pere Molins

Bristol Robotics Laboratory, University of Bristol
Bristol, United Kingdom
molinsp@gmail.com

Sabine Hauert

Bristol Robotics Laboratory, University of Bristol
Bristol, United Kingdom
sabine.hauert@bristol.ac.uk

ABSTRACT

Recent advances in robotics allow for the creation of swarms that work in large numbers. This typically requires low-cost robots with restricted sensing and noisy motion. Developing swarm controllers that are robust to such constraints could lead to applications ranging from outdoor exploration to nanomedicine.

In particular, we demonstrate that a large swarm of minimal robots, using only random motion and limited sensing capabilities, can form trails from a source to an area of interest using a mechanism inspired by diffusion-limited aggregation (DLA). We further show that the nearest area of interest is selected, and that the formed trails can avoid bypass. Validation is performed in simulation and reality using a swarm of up to 100 robots.

CCS CONCEPTS

• **Computing methodologies** → **Cooperation and coordination**;

KEYWORDS

Swarm robotics, minimal robotics, chain formation, path formation, reaction-diffusion systems, diffusion-limited aggregation

1 INTRODUCTION

Large-scale (10^2 and higher) swarms of simple robots are becoming a reality with applications ranging from environmental monitoring to nanomedicine [2, 8]. Robots that work in such large numbers are typically required to be cheap, and as a result minimal in terms of sensing and actuation [17]. New swarm algorithms are needed that are robust to noise and limited robot capabilities [5].

Specifically we look at the formation of trails between a source and an area of interest using minimal robots. Such a trail could be used to guide a crowd to an exit in an emergency situation, or lead a search team to an area of interest, for example a chemical plume, in a large outdoor area [7].

The proposed strategy takes inspiration from a diffusive growth model, namely diffusion-limited aggregation (DLA) [20], to create trails from a source to an area of interest. DLA has successfully been used by [13] in simulation to create trees formed by robot swarms leading to a charging station. Our work modifies DLA to

only create one trail to a single area of interest rather than adopt a tree-like structure between multiple points. We evaluate our control mechanism in simulation and verify the qualitative properties using a swarm of up to 100 real robots. We further show that trails are formed to the nearest area of interest when two are present, and that they are able to bypass obstacles. Rather than improve the performance of trail formation in comparison to other algorithms, our focus is on designing an algorithm that works with minimal robots.

Trail formation has been a widely studied area of swarm robotics. Most work in trail formation takes inspiration from the ability of ants to create trails when foraging for food in nature [3]. Work by Vaughan et al. studied the foraging behavior under the assumption that robots have global positioning and global communication [18]. Systems that only use local information may rely on external beacons. Beacons can either be existing entities in the environment, as is the case with the pre-deployed sensors demonstrated by O'Hara et al. [15], or be deployed dynamically by the robots [1]. Beacons can also be used for indirect, so called stigmergic, communication mimicking the use of pheromones by ants through the storage of virtual pheromone information on the beacons [12].

Figure 1: Up to 100 Kilobot robots were used for swarm experiments.

Another approach relies on the robots acting as the beacons themselves instead of relying on other types of agents or sensors for communication. This can be achieved with robots dynamically switching between beacon and exploration behaviors [10], or through the use of short-range communication to transmit a virtual pheromone gradient [16]. Implementations of the former include the formation of robotic chains for target localization [14, 19] or to order tasks [4], and the creation of communication networks using flying robots [9].

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SAC '19, April 8–12, 2019, Limassol, Cyprus

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5933-7/19/04.

<https://doi.org/10.1145/3297280.3297374>

Most described systems require robots to move in a directed fashion, for example to follow pheromone trails. But deploying large numbers of robots that can perform foraging tasks in large-scale environments, in a cost effective manner, may require further constraints on the agent capabilities. It may be difficult to equip every robot with a GPS for example, or calibrate their motion precisely. Indeed, searching a large area, say a forest, may require tens of thousands of tiny safe-by-design robots. Such scales may benefit from inspiration from microsystems in nature, which typically rely on diffusion and reaction of millions of molecules or cells to interact over proportionally large areas [8].

2 TRAIL FORMATION ALGORITHM

Figure 2: Structures formed by DLA with a seed agent at the center of the environment. Image by Paul Bourke.¹

Diffusion Limited Aggregation (DLA) is a process by which agents aggregate, forming structures that resemble tree branches [11]. A seed agent is located at the center of a circular environment (or a sphere in a three dimensional case) and other agents are released from the perimeter of the circle. These agents start an unbiased random walk until they either leave the environment or encounter the seed, in which case they stick to it. After a few iterations, clusters of agents like the ones shown in Figure 2 will begin to form.

The mechanism proposed in this paper takes inspiration from the aggregation patterns of DLA to create a distributed control strategy capable of creating trails from a source to an area of interest without relying on directional control. We demonstrate that the stochastic processes in systems like DLA can be governed through aggregation parameters, and experiments designed in a way that can be useful to locate areas of interest, even when the agents have extremely limited capabilities.

We define both a source and an area of interest within a circular open environment. The source is placed at the center of the circular environment and the area of interest at a fixed distance. The area of interest emits a signal within a limited range $r_{interest}$. New robotic agents are released at a constant rate ($rate = \frac{\#robots}{time}$) from a source in the center of the environment. Agents start moving randomly at

¹<http://paulbourke.net/fractals/dla/>

Figure 3: Snapshots of a simulation on a circular environment with a source in the middle, and an area of interest located within the environment. The circle on the top shows the environment before the simulation starts. The image in the middle corresponds to a snapshot taken during runtime, and the bottom image represents a snapshot taken at the end of the simulation. A trail can be seen growing from the area of interest to the source.

a constant speed s through the environment. If one of them finds itself within the transmission range $r_{interest}$ of a signal, it will stop moving (from here on referred as binding) and start emitting a signal to robots within communication range $r_{communication}$ for a limited amount of time t . Likewise, mobile robots that enter the signal range of a robot also immobilise and emit a signal for a limited amount of time.

The pseudocode for an implementation of the described mechanism can be found in algorithm 1.

```

if not bound then
  if sense signal then
    | set state to bound;
  else
    | perform random walk;
  end
else
  stop motion;
  if time smaller than maximum signal time then
    | emit signal;
  end
end

```

Algorithm 1: Pseudocode for the trail formation algorithm.

3 SIMULATION

We model idealised robots as particles using GAMA², a platform for building spatially explicit multi-agent simulations designed by Grignard et al. [6]. The advantage of this approach is that we can easily test large-scale systems on the order of 10^3 agents. Dimensionless parameters are used in these simulation, as the goal is to improve our overall understanding of the proposed algorithm, and not to design an algorithm for a specific robot. Parameters used for the simulated scenarios, unless stated otherwise, are $rate=1$, $r_{interest}=20$, $s=65$, $r_{communication}=20$, $t=50$. Simulations are run until a trail is formed, or up to maximum duration of 1000 timesteps. For repeatability, the software and parameters used are provided at <http://hauertlab.com/software>.

The intuition for how the trail is formed is as follows: Figure 4 shows how the diffusion time (defined as the number of timesteps the agents are in a diffusion state and not bound to other agents or the area of interest) of the agents becomes smaller as agents are added at the source, i.e. agents released later during the simulation will spend less time diffusing on average. This is due to the fact that with the continuous release of new agents from the source, the number of agents that bind closer to the source is higher than the number of agents binding further away. The agents emitting signal further away from the source will eventually stop emitting signal, which will further increase the bias for binding close to the source. This leads to the clustering of agents and the formation of the trail visible in Figure 3. When the trail eventually reaches the source, the high-number of robots emitting a signal there will cause any new agent released to bind immediately, effectively blocking the diffusion of newly generated agents. This blocking of the source, together with the fact that agents only emit signal for a short period of time, avoids the further development of new branches.

²<https://github.com/gama-platform>

Figure 4: Average number of timesteps (over 5 runs for the same scenario) taken by agents to bind, as a function of their release order from the source. Agents released later bind quicker on average.

Figure 5: Average distance from the source (over 5 runs for the same scenario) at which agents bind, as a function of their release order from the source. Agents released later bind nearer to the source on average.

Figure 5 shows the binding distance from the source averaged over five runs with the setup from Figure 3. Again, we observe that as time passes, and more agents are released, the agents start binding closer to the source until they block the source. The distance then remains flat to the end of the simulation.

Figure 6 shows the impact of the release rate on trail formation. In particular, we increased the release rate by a factor of 30 and the communication radius by factor of 1.3. Resulting simulations show a more disperse distribution of agents in the environment and no visible trail between the source and the area of interest. We computed the root-mean-square error from a case with ideal performance, a straight line from the source to the area of interest. With the new parameters, the root-mean-square error is 204 as opposed to 169 in the original simulation. The performance decrease is due to the high density of agents present at the time the first agents start binding. This allows for agents to rapidly bind to neighbors in all directions without forming directional patterns. This suggests that a key aspect to obtain efficient and highly directional trails is

Figure 6: Impact of the release rate and communication range on trail formation. Scenarios with higher release rates and larger communication ranges (A) form worse defined trails than scenarios with lower release rates and communication ranges (B).

Figure 7: Final snapshot of a simulation with two areas of interest, with the area of interest on the right placed closer to the source. The trail is formed to the nearest area of interest.

Figure 8: Final snapshot of a simulation with an obstacle between the source and the area of interest. The trail is successfully formed to the source while bypassing the obstacle.

to control the release of agents and therefore the density of agents at the time of the first binding event.

We looked at additional scenarios to test the robustness of the demonstrated results and explore emergent features of the system. The first is a scenario with two areas of interest in the environment, one placed slightly closer to the source than the other. As we can see in Figure 7, the trail will start growing from each area of interest following the described mechanism. Because one of the trails is closer to the source, the probability of one agent binding to it will be higher and thus more agents will bind around that area of interest. Even if agents start aggregating around both areas of interest, the trail from the area of interest which is placed closer to the source will grow faster until it reaches and blocks the source. This will prevent the alternative trail starting at the opposite area of interest from growing, as depicted in Figure 7. This property allows the system to identify the closest area of interest without explicitly obtaining any information about the distance. We also looked at simulations where the two trails are situated at the same distance from the source. In those cases, there is the possibility of creating two trails, but small differences in the amount of agents initially clustering around one of the areas of interest can lead to one of the trails not being able to completely reach the source before the release of agents at the source is blocked by the trail stemming from the other area of interest.

We also explored the robustness of the system by placing an obstacle between the source and the area of interest. In Figure 8, we observe how the trail grows around the obstacle from the area of interest back to the source. To obtain this behavior we increased the speed of the agents by a factor of 1.5.

4 ROBOTIC VALIDATION

We validated the results obtained in simulation using a robotic swarm (Figure 9). This validation is important because much of the cited work [1, 10, 13, 14, 19] is based either exclusively on simulations or on experiments with a small number of robots. We used the Kilobot [17], a low-cost robotic platform designed to scale to many-robot experiments. The Kilobot (shown in Figure 1) is small ($\approx 3\text{cm}$ in diameter) and moves with a maximum speed of 1 cm/s using two vibrating motors that allow it to turn and move forward. It can also communicate omnidirectionally with nearby

Figure 9: Trail formation at the end of the simulation (upper left), and at the end of the robot experiment for three trials. In all cases, trails are formed to the area of interest.

Figure 10: Trail formation at the end of the simulation (upper left), and at the end of the robot experiment for three trials. Two areas of interest are present in the environment, with one of them placed closer to the source. In all cases, trails are formed to the nearest area of interest.

Figure 11: Trail formation at the end of the simulation (upper left), and at the end of the experiment for three trials, when an obstacle is placed along a direct line between the source and the area of interest. In all cases we can observe the adaptation of the trail to circumvent the obstacle.

robots within 10cm using infrared light. Once programmed, the robots are fully autonomous and do not rely on any external input to operate. Experiments are performed on a 2m x 3m arena with videos being captured through an overhead camera. Experiments are stopped as soon as a trail has been formed.

To implement the rules described in algorithm 1 with the robotic swarm, we need to make the robots diffuse in the environment. We programmed each robot to randomly choose a turning direction (left or right) with equal probability and then turn for a random amount of seconds before going straight for another random time interval. The robots are not identical and have differences in turning speed and a direction bias when they go straight (it is not possible to calibrate the two vibrating motors for highly precise straight motion). These inaccuracies are present on an individual level and do not cause a global bias on the entire swarm. The release rate of the agents used in simulation has been replicated by manually placing three robots every 2.5 min. Robots stop being released when the trail has formed to the source, which typically takes more than one hour. The precision of this method is also constrained by human error while timing the intervals. The goal of the experiments was not to obtain highly precise experimental conditions, but rather observe if the overall behavior and qualitative properties of the system showed in simulation could be replicated despite all the experimental constraints. Applications in real world systems will likely involve similar or greater constraints.

We performed the experiments with the robots three times for each of the scenarios. In Figure 9 and Figure 11, we can see a qualitative comparison of the results between the robot experiments

and the simulations. The robotic experiments showed that the three main properties observed in simulation can be replicated with the robots.

The first property shown in Figure 9 (top) is the ability of the system to form trails. We can also observe the emergence of secondary branches among the trail, as in simulation. A video of this behaviour can be found at <https://youtu.be/qZXnq8wYCoU>.

The second property is the capacity to find the closest area of interest without actively sensing any information related to distance. Figure 9(bottom) shows the robots forming a trail towards the area of interest, which is situated closer to the source. This result is also consistent with the results obtained in simulation.

Finally, the ability to form a trail around an obstacle was tested with the robots, as shown in Figure 11 (bottom). As in simulation, the robots are able to grow trails from the area of interest to the source while bypassing the obstacle. The trail further seems to closely follow the contour of the obstacle.

5 CONCLUSIONS AND FUTURE WORK

In this paper we presented a swarm strategy for simple robots that only relies on diffusion and reaction to form trails between two points. In addition to trail formation, the strategy seems to favor the closest area of interest in the environment, and adapts to obstacles. The simplicity of the approach lends itself to swarm deployments in large numbers for applications ranging from exploring unknown environments, to nanomedicine. Future work will focus on a thorough exploration of the parameter space of

the presented algorithm to optimise speed and robustness of trail formation.

REFERENCES

- [1] E.J. Barth. 2003. A dynamic programming approach to robotic swarm navigation using relay markers. *Proceedings of the American Control Conference* 6 (2003), 5264–5269.
- [2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. 2013. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence* 7, 1 (jan 2013), 1–41.
- [3] N. F. Britton, T. R. Stickland, and N. R. Franks. 1998. Analysis of ant foraging algorithms. *Journal of Biological Systems* (1998).
- [4] L. Garattoni and M. Birattari. 2018. Autonomous task sequencing in a robot swarm. *Science Robotics* 3, 20 (2018).
- [5] M. Gauci, J. Chen, W. Li, T. J. Dodd, and R. Groß. 2014. Clustering objects with robots that do not compute. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 421–428.
- [6] A. Grignard, P. Taillandier, B. Gaudou, D. A. Vo, N. Q. Huynh, and A. Drogoul. 2013. GAMA 1.6: Advancing the art of complex agent-based modeling and simulation. In *PRIMA 2013: Principles and Practice of Multi-Agent Systems*, Guido Boella, Edith Elkind, Bastin Tony Roy Savarimuthu, Frank Dignum, and Martin K. Purvis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 117–131.
- [7] H. Hamann. 2018. *Swarm Robotics: A Formal Approach*.
- [8] S. Hauert and S. N. Bhatia. 2014. Mechanisms of cooperation in cancer nanomedicine: towards systems nanotechnology. *Trends in Biotechnology* 32, 9 (2014), 448–455.
- [9] S. Hauert, L. Winkler, J.-C. Zufferey, and D. Floreano. 2008. Ant-based swarming with positionless micro air vehicles for communication relay. *Swarm Intelligence* 2, 2-4 (2008), 167–188.
- [10] Nicholas R. Hoff, Amelia Sagoff, Robert J. Wood, and Radhika Nagpal. 2010. Two foraging algorithms for robot swarms using only local communication. *Proceedings of the IEEE International Conference on Robotics and Biomimetics, ROBIO 2010* (2010), 123–130.
- [11] K. Kassner. 1996. *Pattern formation in diffusion-limited crystal growth: Beyond the single dendrite*. World Scientific Publishing Co Pte Ltd.
- [12] M. Mamei and F. Zambonelli. 2005. Physical deployment of digital pheromones through RFID technology. *Proceedings of the IEEE Swarm Intelligence Symposium (SIS) 2005* (2005), 289–296.
- [13] J. Mullins, B. Meyer, and Aiguo P. Hu. 2012. Collective robot navigation using diffusion limited aggregation. In *Parallel Problem Solving from Nature - PPSN XII*, Carlos A. C. Coello, Vincenzo Cutello, Kalyanmoy Deb, Stephanie Forrest, Giuseppe Nicosia, and Mario Pavone (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 266–276.
- [14] S. Nouyan, A.e Campo, and M. Dorigo. 2008. Path formation in a robot swarm: Self-organized strategies to find your way home. *Swarm Intelligence* 2, 1 (2008), 1–23.
- [15] K. J. O'Hara, D. B. Walker, and T. R. Balch. 2005. *Multi-robot systems. From swarms to intelligent automata*. Springer Netherlands, Dordrecht, Chapter The Gnats – Low-Cost Embedded Networks for Supporting Mobile Robots, 277–282.
- [16] D. W. Payton, M. J. Daily, B. Hoff, M. D. Howard, and C. L. Lee. 2001. Pheromone robotics. In *Intelligent Systems and Smart Manufacturing*. International Society for Optics and Photonics, 67–75.
- [17] M. Rubenstein, C. Ahler, and R. Nagpal. 2012. Kilobot: A low cost scalable robot system for collective behaviors. *Proceedings of the IEEE International Conference on Robotics and Automation* (2012), 3293–3298.
- [18] R. T. Vaughan, K. Støy, G. S. Sukhatme, and M. J. Matarić. 2002. LOST: Localization-space trails for robot teams. *IEEE Transactions on Robotics and Automation* 18, 5 (2002), 796–812.
- [19] B. B. Werger and M. J. Mataric. 1996. Robotic “food” chains: Externalization of state and program for minimal-agent foraging. *From Animals to Animats 4. Proceedings of the International Conference on Simulation of Adaptive Behavior (SAB96)* (1996), 625–634.
- [20] T. A. Witten and L. M. Sander. 1981. Diffusion-limited aggregation, a kinetic critical phenomenon. *Physical Review Letters* 47 (Nov 1981), 1400–1403. Issue 19.