Online Dispersion Algorithms for Swarms of Robots*

Tien-Ruey Hsiang† Esther M. Arkin† Michael A. Bender ‡ Sándor P. Fekete§
Joseph S. B. Mitchell†

1 Introduction

We illustrate algorithms for dispersing a swarm of primitive robots in a two-dimensional unknown environment R. Each robot in the swarm is equipped with a very simple sensor that is able to query the contents of neighboring locations to test the presence of other robots or obstacles. Based on the sensor readings, at each time step a robot may decide to take a discrete step to a neighboring (grid) point. The objective is to minimize the *makespan*, that is, the time to *fill* R with robots, one per grid point inside R.

We focus here on the case of a discrete environment, composed of unit squares (*pixels*) that are induced by the integer grid within a polygonal domain R. There is at most one robot per pixel and robots move horizontally or vertically at unit speed. Robots enter R by means of $k \geq 1$ *door pixels* on the boundary of R, each of which acts as an infinite source of robots.

Robots are primitive finite automata, only having local communication, local sensors, and a constant-sized memory. These local autonomous agents are not centrally controlled, yet are expected to perform the global task of dispersing.

In a recent paper [2], we provide a variety of theoretical results, including:

- Algorithms for the single-door case (i.e., $k = 1$), based on leader-follower strategies that are carefully adapted from the principles of depth-first and breadth-first search to make them applicable to the local nature of the robot model: We prove that our algorithms have optimal makespan $2A - 1$, where A is the area (the number of pixels) of R.
- An algorithm for the multi-door case ($k \geq 1$), based on a wall-following version of the leader-follower strategy: We prove that our strategy is $O(\log(k+1))$-competitive, and that this bound is tight for our strategy as well as other natural "simple-minded" strategies that use only strictly local information.

The purpose of this video is to illustrate via animations the results of implementing our dispersion strategies within a simulator. This has been developed as part of a joint project on "pheromone robotics" for world-embedded computations, under DARPA funding, with HRL Laboratories (David Payton et al. [3]). The dispersion problem studied here has been posed to us by Doug Gage, one of the pioneers of swarm robotics [1].

The video shows how even simple local rules lead to complex emergent behaviors of a swarm of robots. These global behaviors are relatively easy to understand visually, while being difficult to convey via traditional text in a paper.

While our results here are given purely in terms of robots moving synchronously on discrete grids, the results apply also to dispersing swarms of robots within an arbitrary connected planar domain, in the ideal setting in which the robots have perfect motion control and synchrony.

2 The Model

We assume that at most one robot can occupy any one pixel at any given time t. Time is discretized into steps $t = 0, 1, 2, 3, \ldots$. At time 0, there is one robot in each door. Each pixel is in one of three states at time t: (1) the pixel is an *obstacle* if it does not lie in R; (2) the pixel is *occupied* if it is in R and a robot occupies it; or (3) the pixel is *unoccupied* if it is in R and no robot occupies it at time t. An unoccupied pixel at time t is classified as either *previously occupied*, if it was occupied by some robot at some time prior to

*Partially supported by HRL Labs (DARPA subcontract), Honda, NASA Ames Research, NSF (CCR-0098172, EIA-0112849, CCR-0208670), U.S.-Israel Binational Science Foundation, and Sandia National Labs.

†Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY 11794-3600, USA. email: {trhsiang, estie, jsbm} @ams.sunysb.edu.

‡Department of Computer Science, Stony Brook University, Stony Brook, NY 11794-4400, USA. email: bender@cs.sunysb.edu.

§Abt. für Mathematische Optimierung, TU Braunschweig, 38106 Braunschweig, Germany, s.fekete@tu-bs.de.

t, or *frontier*, if it has never been occupied. Robots have *sensors* that detect information about the local environment. In particular, if a robot occupies pixel (i, j) at time t, then we assume that it can detect the state of each pixel within a distance r_S of (i, j), where r_S is the (fixed) *sensor radius*. Robots have no global sensor (e.g., a GPS) and no knowledge of the coordinates of the pixels they occupy.

Each robot has a small finite memory that allows it to remember the sensor readings of the last $T \geq 1$ time steps; i.e., a robot knows the sensor readings it has taken at time $t - T$, $t - T + 1, \ldots, t - 1, t$. In particular, each robot knows which nearby pixels have been occupied recently by other robots.

Each robot has a limited ability to communicate with nearby robots. In particular, at any given time step a robot is able to exchange a constant-size message with any robot that lies within a prescribed *communication radius*, r_C, of the robot. The *communication graph*, $\mathcal{G}(t)$, of the swarm $S(t)$ at time t is defined to be the undirected graph whose nodes are the robots $S(t)$ and whose edges link pairs of robots that communicate. The swarm is said to be *connected* at time t if each connected component of $\mathcal{G}(t)$ contains at least one door pixel.

A *strategy* is a set of rules by which robots move, basing their movement decisions solely on the constant amount of information they sense and remember. The *makespan* of a strategy is the minimum time, t^*, until the robots have filled the region R.

3 The Video Sequence

The video opens with an animation of filling of a complex room. Then the video explains the component of the filling algorithm. The video displays a polygonal room with two doors, each of width one, and shows how the chains of robots fill the domain. Then the video demonstrates the movement rules within each chain, where each robot has to wait until the spot in front of it is vacant before moving.

Next, an animation shows dispersal into a room through a single door ($k = 1$) via a simple depth-first-search strategy. This is followed by a multiple-door example that illustrates the possibility of unwanted "blocking": The chains of robots on the left reach a dead end and are stopped, leaving only the rightmost chain of robots still active to continue to fill the region; this is wasteful because the strategy fails to use the full width of the door.

A solution to the problem of blocking is splicing, where one chain of robots "cuts into" a neighboring chain of robots in order to keep door pixels actively feeding robots into the room. Splicing must be implemented carefully; the video demonstrates the danger of splicing, which can lead to cycling, where robots continually move without filling new territory.

An example is then shown to illustrate the potential problem associated with splicing, allowing a "flow" of robots in one chain to "cut into" another flow of robots: this can readily lead to cycling.

The next short animation shows how cycles can arise from a straightforward leader-follower algorithm, in which a robot that has no frontier pixel next to it chooses to go to any unoccupied neighboring pixel (splicing into the leader-follower flow passing through this pixel).

Our strategy for multiple doors ($k > 1$) avoids the pitfalls of splicing by allowing splicing to occur in a controlled way: a robot can "cut into" a flow only from one side of the flow (the left side), as illustrated in the next short animation sequence. This, coupled with a prioritized movement (that a robot prefers to go left, then to go forward, then to go right), results in a movement strategy that effectively makes the robots "hug" the left wall as much as possible. The robots executing this strategy fill the room in a "laminar flow" behavior, which results in a provable online strategy, with makespan within a factor $O(\log(k+1))$ of optimal [2]. (A figure shown in the video illustrates a construction showing that this competitive ratio is optimal, assuming a strategy using only local information.)

The next animation shows the simulator with an environment having 3 doors (of widths 2, 3, and 7), with the Laminar Flow Leader-Follower (LFLF) algorithm. The video concludes with animations showing the filling of two classes of "combs" whose teeth have widths that grow exponentially: in one case, the left hand on wall LFLF strategy works optimally; in the other case, the strategy requires makespan that is a factor $\Theta(\log k)$ times optimal.

References

[1] D. Gage. Sensor abstractions to support many-robot systems. In *Proc. of SPIE Mobile Robots VII*, Boston, Vol. 1831, pages 235–246, 1992.

[2] T.-R. Hsiang, E. M. Arkin, M. Bender, S. P. Fekete, and J. S. B. Mitchell. Algorithms for rapidly dispersing robot swarms in unknown environments. *5th International Workshop on Algorithmic Foundations of Robotics*, 2002.

[3] D. Payton, M. Daily, R. Estkowski, M. Howard, C. Lee. Pheromone Robotics. *Autonomous Robots*, Vol. 11, No. 3, pages 319-324, 2001.