Targeted Drug Delivery: Algorithmic Methods for Collecting a Swarm of Particles with Uniform, External Forces

Aaron T. Becker[1], Sándor P. Fekete[2], Li Huang[1], Phillip Keldenich[2], Linda Kleist[2], Dominik Krupke[2], Christian Rieck[2], and Arne Schmidt[2]

1 **Department of Electrical and Computer Engineering, University of Houston, USA.** {atbecker,lhuang28}@uh.edu

2 **Department of Computer Science, TU Braunschweig, Germany.** {s.fekete, p.keldenich, l.kleist, d.krupke, c.rieck, arne.schmidt}@tu-bs.de

Abstract

We investigate algorithmic approaches for targeted drug delivery in a complex, maze-like environment, such as a vascular system. The basic scenario is given by a large swarm of micro-scale particles ("agents") and a particular target region ("tumor") within a system of passageways. Agents are too small to contain on-board power or computation and are instead controlled by a global external force that acts uniformly on all particles, such as an applied fluidic flow or electric field. The challenge is to deliver all agents to the target region with a minimum number of actuation steps. We provide a number of results for this challenge. We show that the underlying problem is NP-hard, which explains why previous work did not provide provably efficient algorithms. We also develop a number of algorithmic approaches that greatly improve the worst-case guarantees for the number of required actuation steps.

1 Introduction

A crucial challenge for a wide range of vital medical problems, such as the treatment of cancer, localized infections and inflammation, or internal bleeding is to deliver active substances to a specific location in an organism. The traditional approach of administering a sufficiently large supply of these substances into the circulating blood may cause serious side effects, as the outcome intended for the target site may also occur in other places, with often undesired, serious consequences. Moreover, novel custom-made substances that are specifically designed for precise effects are usually in too short supply to be generously poured into the blood stream. In the context of targeting brain tumors (see Figure 1), an additional difficulty is the blood-brain barrier. This makes it necessary to develop other, more focused methods for delivering agents to specific target regions.

Given the main scenario of medical applications, this requires dealing with navigation through complex vascular systems, in which access to a target location is provided by pathways (in the form of blood vessels) through a maze of obstacles. However, the microscopic size of particles necessary for passage through these vessels makes it prohibitively difficult to store sufficient energy in suitably sized microrobots, in particular in the presence of flowing blood.

A promising alternative is offered by employing a global external force, e.g., a fluidic flow or an electromagnetic field. When such a force is applied, all particles are subjected to the same direction and distance of motion, unless they are blocked by obstacles in their way. While this makes it possible to move all particles at once, it introduces the difficulty of using *uniform* forces for many particles in *different* locations with different local topology to

36th European Workshop on Computational Geometry, Würzburg, Germany, March 16–18, 2020.
This is an extended abstract of a presentation given at EuroCG'20. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

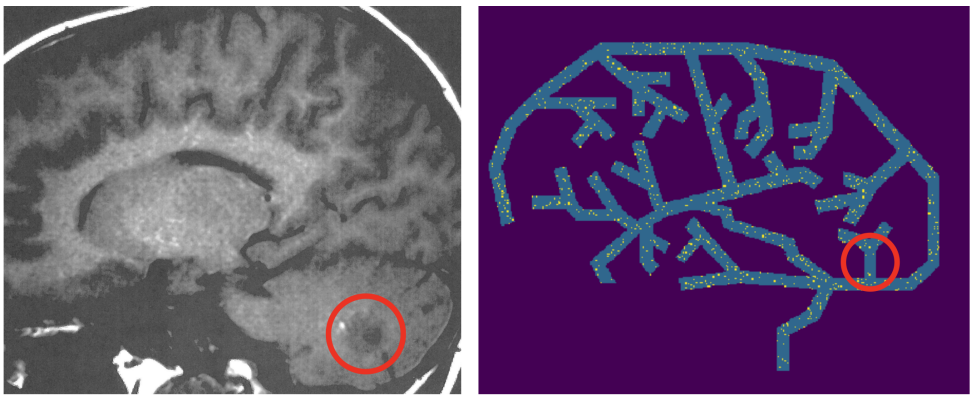

Figure 1 (Left) An MRI image of a brain tumor (red circle), located in the cerebellum. (Right) How can the swarm of particles (yellow dots) be delivered to the target region?

navigate them to *one* final destination. In this paper, we investigate how this objective can be achieved with a small number of actuator steps.

Previous work [11] described a basic approach that delivers all particles in a polyomino with n pixels to a target in at most $O(n^3)$ actuator steps. While a delivery time of this magnitude is usually impractical, we investigate how to improve this.

Our Contribution.

- We prove that minimizing the length of a command sequence for gathering all particles is NP-hard, even if the environments are modeled by polyominoes. Our reduction implies hardness for the related localization problem (as explained in Section 3 before Corollary 3.2).
- We develop an algorithmic strategy for gathering all particles in a polyomino with a worst-case guarantee of at most $O(kD^2)$ steps; here D denotes the maximum distance between any two pixels of the polyomino and k the number of its convex corners. Both k and D are usually much smaller than the number n of grid locations in the polyomino: n may be in $\Omega(D^2)$, for two-dimensional and in $\Omega(D^3)$ for three-dimensional environments.
- For the special case of hole-free polyominoes, we can gather all particles in $O(kD)$ steps.

Further details and algorithmic studies can be found in [7].

1.1 Related Work

This paper seeks to understand control for large numbers of microrobots, and uses a generalized model that could apply to a variety of drug-carrying microparticles. An example are particles with a magnetic core and a catalytic surface for carrying medicinal payloads [10, 15]. An alternative are aggregates of *superparamagnetic iron oxide microparticles*, 9 μm particles that are used as a contrast agent in MRI studies [14]. Real-time MRI scanning can allow feedback control using the location of a swarm of these particles.

Steering magnetic particles using the magnetic gradient coils in an MRI scanner was implemented in [12, 15]. 3D Maxwell-Helmholtz coils are often used for precise magnetic field control [14]. Still needed are motion planning algorithms to guide the swarms of robots through vascular networks. To this end, we build on the techniques for controlling many simple robots with uniform control inputs presented in [4–6]; see video and abstract [3] for a visualizing overview. For a recent survey on challenges related to controlling multiple microrobots (less than 64 robots at a time), see [8].

As the underlying problem consists of bringing together a number of agents in one location, a highly relevant algorithmic line of research considers *rendezvous search*, which requires two or more independent, intelligent agents to meet [1, 2, 9, 13].

2 Preliminaries

The "robots" in this paper are simple particles without autonomy. Every environment is modeled by a *polyomino*, i.e., a set of unit squares, so called *pixels*, in the plane which are joined edge to edge. An example of a polyomino is illustrated in Figure 2. Pixels in the plane not belonging to P are *blocked* because they stop the motion from an adjacent pixel. The particles are commanded in unison: In each step, all particles are relocated by one unit in one of the directions "Up" (u), "Down" (d), "Left" (l), or "Right" (r), unless the destination is a blocked pixel; in this case, a particle remains in its previous pixel. A motion plan is a command sequence $C = \langle c_1, c_2, c_3, \dots \rangle$, where each command $c_i \in \{u, d, l, r\}$.

We assume that the size of a particle is insignificant compared to a pixel. Hence, many of them can be located in the same pixel. During the course of a command sequence, two particles π_1 and π_2 may end up in the same pixel p, if π_1 moves into p, while π_2 remains in p due to a blocked pixel. Once two particles share a pixel, any subsequent command will relocate them in unison—they will not be separated, so they can be considered to be *merged*.

The *distance* between two pixels p and q is the length of a shortest path on the integer grid between p and q that stays within P. The *diameter* of a polyomino P describes the maximum distance between any two of its pixels; we denote it by D. A *configuration* of P is a set of pixels containing at least one particle. The set of all possible configurations of P is denoted by $\mathcal{P}$. We call a command sequence *gathering* if it transforms a configuration $A \in \mathcal{P}$ into a configuration A' such that $|A'| = 1$, i.e., if it merges all particles in the same pixel.

3 Hardness

We show that the following decision problem, which we call Min-Gathering, is hard: Given a polyomino P and a set of particles, is there a gathering sequence of length ℓ?

▶ **Theorem 3.1.** Min-Gathering *is NP-hard.*

Proof-Sketch. The proof is based on a reduction from 3-Sat. For every instance Φ of 3-Sat, we construct a polyomino P_Φ of diameter D containing a particle in every pixel such that there exists a gathering sequence of length $\ell := \frac{1}{2}(D + b)$ if and only if Φ is satisfiable.

P_Φ is constructed as follows, see Figure 2: For every variable, we insert a variable gadget. We join all variable gadgets vertically in a row to a *variable block*; we call the top row of each variable gadget its *variable row*. For every clause, we construct a clause gadget that contains a left (right) *literal arm* for each incident positive (negative) literal in the corresponding variable row and an exit arm in the bottom. To obtain P_Φ, we join all clause gadgets from left to right by a *bottom row* and insert a variable block at the left and right end of the bottom row of length b. Note that b denotes the number of pixels in the bottom row of P_Φ, and that the distance between the two *red* particles (i.e., the two leftmost particles above the variable blocks) realizes the diameter D.

We can argue that by applying a command sequence according to a satisfying assignment for Φ, the left (right) red particle moves to the left (right) pixel of the bottom row, where these particles can be merged, yielding a gathering sequence of length ℓ. Note that in this command sequence, particles in the clause gadgets traverse one of the literal arms, reaching the bottom row at the exact same time.

Conversely, if Φ is not satisfiable, then in at least one clause a particle uses the exit arm (see blue particle). Therefore, due to the different heights of the exit arms of the variable and the clause gadgets, particles remain in the exit arms and do not reach the bottom row simultaneously with the red particles. Thus, there is no gathering sequence of length ℓ. ◀

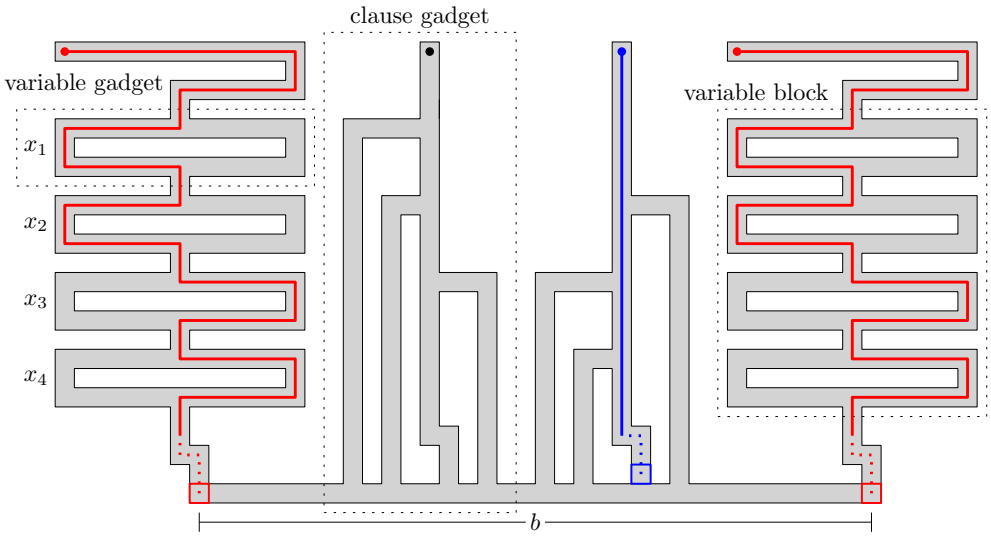

Figure 2 The polyomino P_Φ for the 3-SAT-instance $\Phi = (x_1 \vee x_2 \vee \overline{x}_3) \wedge (\overline{x}_2 \vee x_3 \vee x_4)$. A gathering sequence for the two red particles of length $\frac{1}{2}(D+b)$ corresponds to a variable assignment of Φ.

Note that the left pixel of the bottom row is one of two possible merge location for a gathering sequence of length $\frac{1}{2}(D+b)$. Therefore, the same reduction shows that the problem remains hard if a target location is prescribed. Because every pixel contains a particle, this implies that the decision problem of ROBOT LOCALIZATION is also hard. In an instance of this problem, we are given a sensorless robot r in a polyomino, and wonder whether there exists a command sequence of length ℓ such that we know the position of r afterwards. The above observations yield the following.

▶ **Corollary 3.2.** ROBOT LOCALIZATION *is NP-hard.*

4 Algorithmic Approaches

We start by merging two particles in a special class of polyominoes. A polyomino P is *simple* (or *hole-free*) if decomposing P with horizontal lines through pixel edges results in a set of rectangles $\mathcal{R}$ such that the edge-contact graph $\mathcal{C}(\mathcal{R})$ of $\mathcal{R}$ is a tree. The edge-contact graph of $\mathcal{R}$ has a vertex for each rectangle and an edge for each side contact; a point contact does not suffice.

▶ **Theorem 4.1.** *For any two particles in a simple polyomino P with diameter D, there exists a gathering sequence of length D.*

Proof-Sketch. Let $\mathcal{R}$ be a decomposition of P into rectangles by cutting P with horizontal lines through pixel edges. Then, because P is simple, the edge-contact graph $\mathcal{C}(\mathcal{R})$ of the rectangles $\mathcal{R}$ is a tree. For an example, consider Figure 3.

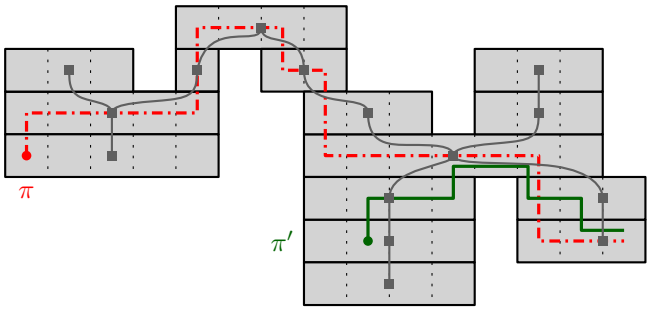

Figure 3 A simple polyomino P, and its edge-contact graph $\mathcal{C}(\mathcal{R})$ (in gray). When the red particle π moves towards the green particle π', π and π' follow the red and the green path, respectively.

For every t, let R_t and R'_t be the rectangles of P containing the two particles π and π' after applying t commands, respectively. Moreover, let S_t be a shortest path from R_t to R'_t in $\mathcal{C}(\mathcal{R})$. Moreover, let $S_t(1)$ be the successor of R_t on S_t (if it exists, i.e., $R_t \neq R'_t$).

We use the following strategy.

Phase 1: While $R_t \neq R'_t$, compute a shortest path S_t from R_t to R'_t in $\mathcal{C}(\mathcal{R})$. Move π to $S_t(1)$ via a shortest path in P. Update R_t and R'_t.

Phase 2: If $R_t = R'_t$, move π towards π' by a shortest (horizontal) path; note that this gathering sequence merges the particles within R_t.

In fact, the resulting sequence has the following property: For every $s > t$, the rectangles R_s and R'_s are either equal to R_t or lie in the connected component C of $\mathcal{C}(\mathcal{R} \setminus R_t)$ containing R'_t. This implies that the merge location and R'_t lie in C or are equal to R_t. Consequently, in every step, π moves towards the merge location on a shortest path and thus the gathering sequence is at most of length D. ◀

We call the strategy used to prove Theorem 4.1 DynamicShortestPath (DSP): Move one particle towards the other along a shortest path; update the shortest path if a shorter one exists. The example in Figure 4 shows that DSP may perform significantly worse in non-simple polyominoes, i.e., it may not yield a gathering sequence of length $O(D)$.

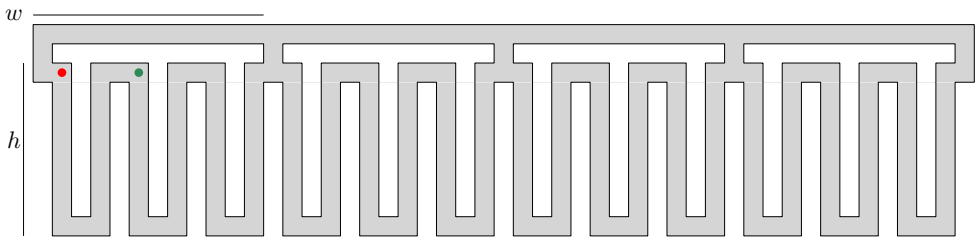

Figure 4 When the red particle π moves towards the green particle π' by shortest paths, π visits the entire bottom path.

Nevertheless, DSP always merges two particles: When a particle π follows π' in a polyomino with n pixels, then within n commands either the shortest path is updated or π' must meet a wall. Therefore, for every n commands, the distance between the particles decreases by at least 1. Therefore, the following holds true.

▶ **Proposition 4.2.** *For every polyomino P with n pixels and diameter D and every configuration with two particles, DSP yields a gathering sequence of length $O(nD)$.*

Using a different strategy yields a better bound.

▶ **Theorem 4.3.** *For any two particles in a polyomino P, there exists a gathering sequence of length at most D^2.*

Proof. Let q be the top-rightmost pixel of P. To merge the two particles in q, our strategy is as follows: Identify the particle π that is bottom-leftmost. Apply a command sequence that moves π to q on a shortest path. Repeat.

▶ **Claim.** *In each iteration, the sum of the distances Δ of the two particles to q decreases.*

Note that Δ decreases when the other particle π' has a collision. If π' had no collision, there exist a pixel that is higher or more to the right than q, contradicting the choice of q. Consequently, the sum of distances Δ, which is at most $2D$ at start, decreases at least by 1 for every D steps. Hence after $O(D^2)$ steps, Δ is reduced to 0. ◀

Note that there exist polyominoes, e.g., a square, where the number n of pixels is in $\Omega(D^2)$. Therefore, Theorem 4.3 significantly improves the bound of $O(n^3)$ in [11]. Finally, we note that a shortest gathering sequence for two particles in a non-simple polyomino may need to exceed D; Figure 5 illustrates the non-simple polyomino used to obtain Proposition 4.4.

▶ **Proposition 4.4.** *There exists a non-simple polyomino P with two particles such that a shortest gathering sequence has length $3/2D - O(\sqrt{D})$.*

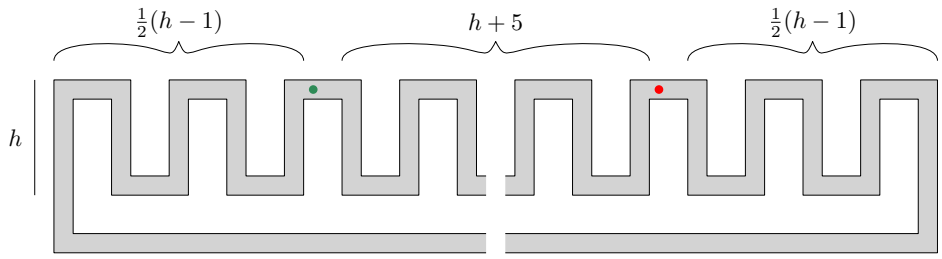

Figure 5 Merging the two particles in this non-simple polyomino needs to exceed D.

In the following, we show how to guarantee with few commands that the number of remaining particles is proportional to the complexity of the polyomino, namely the number of its convex corners.

▶ **Lemma 4.5.** *Let P be a polyomino with diameter D and k convex corners. For every configuration $A \in \mathcal{P}$, there exists a command sequence of length $2D$ which transforms A to a configuration $A' \in \mathcal{P}$ such that $|A'| \leq k/4$.*

Proof. We distinguish four types of convex corners; northwest (NW), northeast (NE), southwest (SW), southeast (SE). By the pigeon hole principle, one of the types occurs at most $k/4$ times; without loss of generality, let this be the NW corners.

We show that after applying the sequence $\langle l, u\rangle^D$, every particle lies in a NW corner: Consider a particle π in pixel p. Unless π lies in a NW corner, it moves for at least one command in $\{l, u\}$. Because P is finite, there exists an ℓ large enough such that π ends in a NW corner q when the command sequence $\langle l, u\rangle^\ell$ is applied, i.e., there exists an pq-path consisting of at most ℓ commands of types l and u, respectively. Because a monotone path is a shortest path, it holds that $\ell \leq D$. ◀

By combining Lemma 4.5 with Theorem 4.1 and Theorem 4.3, respectively, we obtain the following upper bounds.

▶ **Corollary 4.6.** *For a set of particles in a simple polyomino P with diameter D and k convex corners, there exists a gathering sequence of length $O(kD)$.*

▶ **Corollary 4.7.** *For any set of particles in a polyomino P with diameter D and k convex corners, there exists a gathering sequence of length at most $O(kD^2)$.*

References

1. Steve Alpern and Shmuel Gal. *The theory of search games and rendezvous*. International Series in Operations Research and Management Science. Kluwer Academic Publishers, Boston, Dordrecht, London, 2003.
2. Edward J Anderson and Sándor P Fekete. Two dimensional rendezvous search. *Operations Research*, 49(1):107–118, 2001.
3. A. T. Becker, Erik D. Demaine, Sándor P. Fekete, S. H. Mohtasham Shad, and R. Morris-Wright. Tilt: The video. Designing worlds to control robot swarms with only global signals. In *31st International Symposium on Computational Geometry (SoCG)*, pages 16–18, 2015.
4. Aaron T. Becker, Erik D. Demaine, S. P. Fekete, and James McLurkin. Particle computation: Designing worlds to control robot swarms with only global signals. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6751–6756, 2014.
5. Aaron T. Becker, Erik D. Demaine, Sàndor P. Fekete, Golnaz Habibi, and James McLurkin. Reconfiguring massive particle swarms with limited, global control. In *Algorithms for Sensor Systems (ALGOSENSORS)*, pages 51–66, 2014.
6. Aaron T. Becker, Erik D. Demaine, Sándor P. Fekete, Jarrett Lonsford, and Rose Morris-Wright. Particle computation: Complexity, algorithms, and logic. *Natural Computing*, 18(1):181–201, 2019.
7. Aaron T. Becker, Sándor P. Fekete, Li Huang, Phillip Keldenich, Linda Kleist, Dominik Krupke, Christian Rieck, and Arne Schmidt. Targeted Drug Delivery: Algorithmic Methods for Collecting a Swarm of Particles with Uniform, External Forces. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020. To appear.
8. Sagar Chowdhury, Wuming Jing, and David J. Cappelleri. Controlling multiple microrobots: recent progress and future challenges. *Journal of Micro-Bio Robotics*, 10(1-4):1–11, 2015.
9. Paola Flocchini. Gathering. In Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro, editors, *Distributed Computing by Mobile Entities, Current Research in Moving and Computing*, pages 63–82. Springer, 2019.
10. Julia Litvinov, Azeem Nasrullah, Timothy Sherlock, Yi-Ju Wang, Paul Ruchhoeft, and Richard C Willson. High-throughput top-down fabrication of uniform magnetic particles. *PloS one*, 7(5):e37440, 2012.
11. Arun V. Mahadev, Dominik Krupke, Jan-Marc Reinhardt, Sándor P. Fekete, and Aaron T. Becker. Collecting a swarm in a 2D environment using shared, global inputs. In *13th Conference on Automation Science and Engineering (CASE)*, pages 1231–1236, 2016.
12. Jean-Baptiste Mathieu and Sylvain Martel. Magnetic microparticle steering within the constraints of an MRI system: proof of concept of a novel targeting approach. *Biomedical microdevices*, 9(6):801–808, 2007.
13. Malika Meghjani and Gregory Dudek. Multi-robot exploration and rendezvous on graphs. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5270–5276, 2012.
14. Lyes Mellal, David Folio, Karim Belharet, and Antoine Ferreira. Magnetic microbot design framework for antiangiogenic tumor therapy. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1397–1402, 2015.

15 Pierre Pouponneau, Jean-Christophe Leroux, and Sylvain Martel. Magnetic nanoparticles encapsulated into biodegradable microparticles steered with an upgraded magnetic resonance imaging system for tumor chemoembolization. *Biomaterials*, 30(31):6327–6332, 2009.