

# Reconfiguring Massive Particle Swarms with Limited, Global Control

Aaron Becker<sup>1</sup>, Erik D. Demaine<sup>2</sup>, Sándor P. Fekete<sup>3</sup>(✉), Golnaz Habibi<sup>1</sup>,  
and James McLurkin<sup>1</sup>

<sup>1</sup> Department of Computer Science, Rice University, Houston, TX 77005, USA  
`aabecker@gmail.com`, `{gh4,jm23}@rice.edu`

<sup>2</sup> Computer Science and Artificial Intelligence Laboratory, MIT,  
Cambridge, MA 02139, USA  
`edemaine@mit.edu`

<sup>3</sup> Department of Computer Science, TU Braunschweig,  
Mühlenpfordtstr. 23, 38106 Braunschweig, Germany  
`s.fekete@tu-bs.de`

**Abstract.** We investigate algorithmic control of a large swarm of mobile particles (such as robots, sensors, or building material) that move in a 2D workspace using a global input signal such as gravity or a magnetic field. Upon activation of the field, each particle moves maximally in the same direction, until it hits a stationary obstacle or another stationary particle. In an open workspace, this system model is of limited use because it has only two controllable degrees of freedom—all particles receive the same inputs and move uniformly. We show that adding a maze of obstacles to the environment can make the system drastically more complex but also more useful. The resulting model matches ThinkFun’s Tilt puzzle.

If we are given a fixed set of stationary obstacles, we prove that it is NP-hard to decide whether a given initial configuration can be transformed into a desired target configuration. On the positive side, we provide constructive algorithms to design workspaces that efficiently implement arbitrary permutations between different configurations.

**Keywords:** Robot swarm · Nano-particles · Uniform inputs · Parallel motion planning · Complexity · Array permutations

## 1 Introduction

Since the first visions of massive sensor swarms, more than ten years of work on sensor networks have yielded considerable progress with respect to hardware miniaturization. The original visions of “Smart Paint” [1] or “Smart Dust” [27] have triggered a considerable amount of theoretical research on swarms of *stationary* processors, e.g., the work in [16–18, 29]. Recent developments in the ability to design, produce, and control particles at the nanoscale and the rise of possible applications, e.g., targeted drug delivery, micro and nanoscale construction, and Lab-on-a-Chip, motivate the study of large swarms of *mobile* objects.

But how can we control such a swarm with only limited computational power and a lack of individual control by a central authority? Local, robotics-style motion control by the particles themselves appears hopeless, because (1) the physics of motion at the low Reynold’s number nanoscale environment requires overcoming a considerable amount of resistance, and (2) the capacity for storing energy for computation, communication, and motion control shrinks with the third power of object size.

A possible answer lies in applying a global, external force to all particles in the swarm. This resembles the logic puzzle Tilt [38], and dexterity ball-in-a-maze puzzles such as Pigs in Clover and Labyrinth, which involve tilting a board to cause all mobile pieces to roll or slide in a desired direction. Problems of this type are also similar to sliding-block puzzles with fixed obstacles [10, 24–26], except that all particles receive the same control inputs, as in the Tilt puzzle. In the real world, driving ferromagnetic particles with a magnetic resonance imaging (MRI) scanner gives a nano-scale example of this challenge [40]. Becker et al. [7] demonstrate how to apply a magnetic field to simultaneously move cells containing iron particles in a specific direction within a fabricated workspace; see Fig. 1a. Other recent examples include using the global magnetic field from an MRI to guide magneto-tactic bacteria through a vascular network to deliver payloads at specific locations [8], and using electromagnets to steer a magneto-tactic bacterium through a micro-fabricated maze [28]; however, this still involves only individual particles at a time, not the parallel motion of a whole, massive swarm. How can we manipulate the overall swarm with coarse global control, such that individual particles arrive at multiple different destinations in a (known) complex vascular network such as the one in Fig. 1b?


Thus, we study the following basic problem: *Given a map of an environment, such as the vascular network shown in Fig. 1b, along with initial and goal positions for each particle, does there exist a sequence of inputs that will bring each particle to its goal position?*

As it turns out, the deliberate use of existing stationary obstacles leads to a wide range of possible sequences of moves. In the first part of the paper, we show that this may lead to computationally difficult situations. In the second part of the paper (Sect. 5), we develop several positive results. The underlying idea is to construct artificial obstacles (such as walls) that allow arbitrary rearrangements of a given two-dimensional particle swarm. For clearer notation, we will formulate the relevant statements in the language of matrix operations, which is easily translated into plain geometric language.


Our paper is organized as follows. After a formal problem definition in Sect. 2 and a discussion of related work in Sect. 3, we provide our main result on the complexity of the problem in Sect. 4. We then present constructive algorithmic results in Sect. 5, and end with concluding remarks in Sect. 6.

## 2 Problem Definition

We study the problem on a two-dimensional grid. We assume that particles cannot be individually controlled, but are all simultaneously given a message



(a) (Left, center) after feeding iron particles to ciliate eukaryon (*Tetrahymena pyriformis*) and magnetizing the particles with a permanent magnet, the cells can be turned by changing the orientation of an external magnetic field. (Right) using two orthogonal Helmholtz electromagnets (left), Becker et al. demonstrated steering many living magnetized *T. pyriformis* cells [7]. All cells are steered by the same global field.



(b) Biological vascular network (cottonwood leaf). Photo: Royce Bair/Flickr/Getty Images. Given such a network along with initial and goal positions of  $N$  particles, is it possible to bring each particle to its goal position using a global control signal? Note that this arrangement is *not* a tree, but is a graph structure with loops. MATLAB code for driving  $n$  robots through this network available at <http://www.mathworks.com/matlabcentral/fileexchange/42892>.

**Fig. 1.** (Top) State of the art in controlling small objects by force fields. (Bottom) A complex vascular network, forming a typical environment for the parallel navigation of small objects.

to travel in a given direction until they collide with an obstacle or another particle. This assumption corresponds to situations with limited state feedback, or for particles that move at unpredictable speeds. More precisely we consider the following scenario, illustrated in Fig. 2, which we call GLOBALCONTROL-MANYPARTICLES:

1. Initially, the planar square grid is filled with some unit-square particles (each occupying a cell of the grid) and some fixed unit-square blocks.
2. All particles are commanded in unison: a valid command is “Go Up” ( $u$ ), “Go Right” ( $r$ ), “Go Down” ( $d$ ), or “Go Left” ( $l$ ). All particles move in the commanded direction until they hit an obstacle or another particle. A representative command sequence is  $\langle u, r, d, l, d, r, u, \dots \rangle$ . We call these global commands *force-field moves*. We assume we can bound the minimum particle speed and can guarantee all particles have moved to their maximum extent.
3. The goal is to get any particle to a specified position.

The algorithmic decision problem GLOBALCONTROL-MANYPARTICLES is to decide whether a given puzzle is solvable. As it turns out, this problem is computationally difficult: we prove NP-hardness in Sect. 4. While this result shows the richness of our model (despite the limited control over the individual parts), it also constitutes a major impediment for constructive algorithmic work.



**Fig. 2.** In this image, black cells are fixed, white cells are free, solid discs are individual particles, and goal positions are dashed circles. For the simple world at left, it is impossible to maneuver both particles to end at their goals. The world at right has a finite solution:  $\langle r, d, l \rangle$ .

### 3 Related Work

*Large Robot Populations.* Due to the efforts of roboticists, biologists, and chemists (e.g. [9, 35, 37]), it is now possible to make and field very large ( $10^3$ – $10^{14}$ ) populations of simple robots. Potential applications for these robots include targeted medical therapy, sensing, and actuation. With large populations come two fundamental challenges: (1) how to perform state estimation for the robots, and (2) how to control these robots.

Traditional approaches often assume independent control signals for each robot, but each additional independent signal requires bandwidth and engineering. These bandwidth requirements grow at  $O(n)$ . Using independent signals becomes more challenging as the robot size decreases. At the molecular

scale, there is a bounded number of modifications that can be made. Especially at the micro- and nano-scales it is not practical to encode autonomy in the robots. Instead, the robots are controlled and interacted with using global control signals.

More recently, robots have been constructed with physical heterogeneity so that they respond differently to a global, broadcast control signal. Examples include *scratch-drive microrobots*, actuated and controlled by a DC voltage signal from a substrate [12]; magnetic structures with different cross-sections that could be independently steered [19]; *MagMite* microrobots with different resonant frequencies and a global magnetic field [20]; and magnetically controlled nanoscale helical screws constructed to stop movement at different cutoff frequencies of a global magnetic field [36]. In our previous work with robots modeled as nonholonomic unicycles, we showed that an inhomogeneity in turning speed is enough to make even an infinite number of robots controllable with regard to position. All these approaches show promise, but they require precise state estimation and heterogeneous robots. In addition, the control law computation required at best a summation over all the robot states  $O(n)$  [6] and at worst a matrix inversion  $O(n^{2.373})$  [4].

In this paper we take a very different approach. We assume a population of approximately identical planar particles (which could be small robots) and one global control signal that contains the direction all particles should move. In an open environment, this system is not controllable because the particles move uniformly—implementing any control signal translates the entire group identically. However, an obstacle-filled workspace allows us to break symmetry. We showed that if we can command the particles to move one unit distance at a time, some goal configurations have easy solutions [5]. Given a large free space, we have an algorithm showing that a single obstacle is sufficient for position control of  $N$  particles (video of position control: [http://youtu.be/5p\\_XIad5-Cw](http://youtu.be/5p_XIad5-Cw)). This result required incremental position control of the group of particles, i.e. the ability to advance them a uniform fixed distance. This is a strong assumption, and one that we relax in this work.

*Dexterity Games.* The problem we investigate is strongly related to dexterity puzzles—games that typically involve a maze and several balls that should be maneuvered to goal positions. Such games have a long history. *Pigs in Clover*, involving steering four balls through 3 concentric incomplete circles, was invented in 1880 by Charles Martin Crandall. Dexterity games are dynamic and depend on the manual skill of the player. Our problem formulation also applies the same input to every agent, but imposes only kinematic restrictions on agents. This is most similar to the gravity-fed logic maze *Tilt*<sup>TM</sup>, invented by Vesa Timonen and Timo Jokitalo and distributed by [ThinkFun](#) since 2010 [38].

*Computational Geometry: Robot Box-Pushing.* Many variations of block-pushing puzzles have been explored from a computational complexity viewpoint, with a seminal paper proving NP-hardness by Gordon Wilfong in 1991 [43]. The general case of motion-planning when each command moves particles a single unit in a



world composed of even a single robot and both *fixed* and *moveable* squares is in the complexity class PSPACE-COMPLETE [11, 13, 25].

Ricochet Robots [14], Atomix [26], and PushPush [10] have the same constraint that robots when moved must move to their full extent. This constraint reflects physical realities where, due to uncertainties in sensing, control application, and robot models, precise quantified movements in a specified direction is not possible, but the input can be applied for a long period of time and be guaranteed that the robots will move to their fullest extent. In these games the robots move to their full extent with each input, but each robot can be actuated individually. The complexity of the problem with global inputs to all robots has remained an open problem.

*Sensorless Manipulation.* The algorithms in the second half of our paper do not require feedback, and we have drawn inspiration from work on sensorless manipulation [15]. The basic idea in this work is to explicitly maintain the set of all possible robot configurations and to select a sequence of actions that reduces the size of this set and drives it toward some goal configuration. Carefully selected primitive operations can make this easier. For example, sensorless manipulation strategies often use a sequential composition of primitive operations, “squeezing” a part either virtually with a programmable force field or simply between two flat, parallel plates [23]. Some sensorless manipulation strategies take advantage of limit cycle behavior, for example engineering fixed points and basins of attraction so that parts only exit a feeder when they reach the correct orientation [31, 33]. These two strategies have been applied to a much wider array of mechanisms such as vibratory bowls and tables [21, 41, 42] or assembly lines [2, 23, 39], and have also been extended to situations with stochastic uncertainty [22, 32] and closed-loop feedback [3, 34].

*Parallel Algorithms: SIMD.* Another related area of research is Single Instruction Multiple Data (SIMD) parallel algorithms [30]. In this model, multiple processors are all fed the same instructions to execute, but they do so on different data. This model has some flexibility, for example allowing command execution selectively only on certain processors and no operations (NOPs) on the remaining processors.

Our model is actually more extreme: the particles all respond in effectively the same way to the same instruction. The only difference is their location, and which obstacles or particles will thus block them. In some sense, our model is essentially Single Instruction, Single Data, Multiple Location.

## 4 Complexity


We prove that the general problem defined in Sect. 2 is computationally intractable:

**Theorem 1.** GLOBALCONTROL-MANYPARTICLES is NP-hard: given an initial configuration of movable particles and fixed obstacles, it is NP-hard to decide whether any particle can be moved to a specified location.

*Proof.* We prove hardness by a reduction from 3SAT. Suppose we are given  $n$  Boolean variables  $x_1, x_2, \dots, x_n$ , and  $m$  disjunctive clauses  $C_j = U_j \vee V_j \vee W_j$ , where each literal  $U_j, V_j, W_j$  is of the form  $x_i$  or  $\neg x_i$ . We construct an instance of GLOBALCONTROL-MANYPARTICLES that has a solution if and only if all clauses can be satisfied by a truth assignment to the variables.

*Variable gadgets.* For each variable  $x_i$  that appears in  $k_i$  literals, we construct  $k_i$  instances of the *variable gadget*  $i$  shown in Fig. 3, with a particle initially at the top of the gadget. The gadget consists of a tower of  $n$  levels, designed for the overall construction to make  $n$  total variable choices. These choices are intended to be made by a move sequence of the form  $\langle d, l/r, d, l/r, \dots, d, l/r, d, l \rangle$ , where the  $i$ th  $l/r$  choice corresponds to setting variable  $x_i$  to either true ( $l$ ) or false ( $r$ ). Thus variable gadget  $i$  ignores all but the  $i$ th choice by making all other levels lead to the same destination via both  $l$  and  $r$ . The  $i$ th level branches into two destinations, chosen by either  $l$  or  $r$ , which correspond to  $x_i$  being set true or false, respectively.

In fact, the command sequence may include multiple  $l$  and  $r$  commands in a row, in which case the last  $l/r$  before a vertical  $u/d$  command specifies the final decision made at that level, and the others can be ignored. The command sequence may also include a  $u$  command, which undoes a  $d$  command if done immediately after, or else does nothing; thus we can simply ignore the  $u$  command and the immediately preceding  $d$  if it exists. We can also ignore duplicate commands (e.g.,  $d, d$  becomes  $d$ ) and remove any initial  $l/r$  command. After ignoring these superfluous commands, assuming a particle reaches one of the output channels, we obtain a sequence in the canonical form  $\langle d, l/r, d, l/r, \dots, d, l \rangle$  as desired, corresponding uniquely to a truth assignment to the  $n$  variables. (If no particle reaches the output port, it is as if the variable is neither true nor false, satisfying no clauses.) Note that all particles arrive at their output ports at exactly the same time.




**Fig. 3.** Variable gadgets that execute by a sequence of  $\langle d, l/r \rangle$  moves. The  $i$ th  $l/r$  choice sets the variable to true or false by putting the ball in a separate column. This selection move is shown in blue. Each gadget is designed to respond to the  $i$ th choice but ignore all others. This lets us make several copies of the same variable by making multiple gadgets with the same  $i$ . In the figure  $n = 4$ , and the input sequence  $\langle d, l, d, r, d, l, d, r, d, r, d \rangle$  causes  $i = (1, 2, 3, 4)$  to produce (true, false, true, false) (Colour figure online).

*Clause gadgets.* For each clause, we use the OR gadget shown in Fig. 4a. The OR gadget has three inputs corresponding to the three literals, and input particles are initially at the top of these inputs. For each positive literal of the form  $x_i$ , we connect the corresponding input to the left output of an unused instance of variable gadget  $i$ . For each negative literal of the form  $\neg x_i$ , we connect the corresponding input to the right output of an unused instance of a variable gadget  $i$ . (In this way, each variable gadget gets used exactly once.)

We connect the variable gadget to the OR gadget in a simple way, as shown in Fig. 5: place the variable gadget above the clause so as to align the vertical output and input channels, and join them into a common channel. To make room for the three variable gadgets, we simply extend the black areas separating the three input channels in the OR gadget. The unused output channel of each variable gadget is connected to a waste receptacle. Any particle reaching that end cannot return to the logic.

If any input channel of the OR gadget has a particle, then it can reach the output port by the move sequence  $\langle d, l, d, r \rangle$ . Furthermore, because variable gadgets place all particles on their output ports at the same time, if more than one particle reaches the OR gadget, they will move in unison as drawn in Fig. 4a, and only one can make it to the output port; the others will be stuck in the “waste” row, even if extra  $\langle l, r, u, d \rangle$  commands are interjected into the intended sequence. Hence, a single particle can reach the output of a clause if and only if that clause (i.e., at least one of its literals) is satisfied by the variable assignment.



**Fig. 4.** Gadgets that use the cycle  $\langle d, l, d, r \rangle$ . The 3-input OR gadget outputs one particle if at least one particle enters in an input line, and sends any extra particle(s) to a waste receptacle. The  $m$ -input AND gadget outputs one particle to the TARGET LOCATION, marked in gray, if at least  $m$  inputs are TRUE. Here  $m = 5$ . Excess particles are sent to a waste receptacle.

*Check gadget.* As the final stage of the computation, we check that all clauses were simultaneously satisfied by the variable assignment, using the  $m$ -input AND gadget shown in Fig. 4b and c. Specifically, we place the clause gadgets along a horizontal line, and connect their vertical output channels to the vertical input channels of the check gadget. Again we can align the channels by extending the black areas that separate the input channels of the AND gadget, as shown in the composite diagram Fig. 5.

The intended solution sequence for the AND gadget is  $\langle d, l, d, r \rangle$ . The AND gadget is designed with the downward channel exactly  $m$  units to the right from






(a) Initial state. The objective is to get one particle to the grey square at lower left.



(b) Setting variables to (False, True, False, True) does not satisfy this 3SAT problem.



(c) Setting the variables (True, False, False, True) will satisfy this 3SAT problem.



(d) Successful outcome. (True, False, False, True) places a single particle in the goal.

**Fig. 5.** Combining 12 variable gadgets, three 3-input OR gadgets, and an  $m$ -input AND gadget to realize the 3SAT expression  $(\neg A \vee \neg C \vee D) \wedge (\neg B \vee \neg C \vee D) \wedge (\neg A \vee B \vee D) \wedge (A \vee \neg B \vee C)$ . MATLAB code implementing the examples for each figure in the paper is available online <http://www.mathworks.com/matlabcentral/fileexchange/42892>.

the left wall, and  $> 2m$  units from the right wall, so for any particle to reach the downward channel (and ultimately, the target location), at least  $m$  particles must be presented as input. Because each input channel will present at most one particle (as argued in a clause), a particle can reach the final destination if and only if all  $m$  clauses output a particle, which is possible exactly when all clauses are satisfied by the variable assignment.

This completes the reduction and the NP-hardness proof.

We conjecture that GLOBALCONTROL-MANYPARTICLES is in fact PSPACE-COMplete. One approach would be to simulate nondeterministic constraint logic [24], perhaps using a unique move sequence of the form  $\langle d, l/r, d, l/r, \dots \rangle$  to identify and “activate” a component. One challenge is that all gadgets must properly reset to their initial state, without permanently trapping any particles. We leave this for future work.

## 5 Matrix Permutations

The previous sections investigated pathologically difficult configurations. This section investigates a complementary problem. Given the same particle and world constraints as before, what types of control are possible and economical if we are free to design the environment?

First, we describe an arrangement of obstacles that implement an arbitrary matrix permutation in four commands. Then we provide efficient algorithms for sorting matrices, and finish with potential applications.

### 5.1 A Workspace for a Single Permutation


For our purposes, a *matrix* is a 2D array of particles (each possibly a different color). For an  $a_r \times a_c$  matrix  $A$  and a  $b_r \times b_c$  matrix  $B$ , of equal total size  $N = a_r \cdot a_c = b_r \cdot b_c$ , a *matrix permutation* assigns each element in  $A$  a unique position in  $B$ . Figures 6 and 7 show example constructions that execute matrix permutations of total size  $N = 25$  and 100, respectively. For simplicity of exposition, we assume henceforth that all matrices are  $n \times n$  squares.

**Theorem 2.** *Any matrix permutation can be executed by a set of obstacles that transforms matrix  $A$  into matrix  $B$  in just four moves. For  $N$  particles, the arrangement requires  $(3N + 1)^2$  space,  $4N + 1$  obstacles, and  $12N/\text{speed}$  time.*


*Proof.* Refer to Figs. 6 and 7 for examples. The move sequence is  $\langle u, r, d, l \rangle$ .

**Move 1:** We place  $n$  obstacles, one for each column, spaced  $n$  units apart, such that moving  $u$  spreads the particle array into a staggered vertical line. Each particle now has its own row. **Move 2:** We place  $N$  obstacles to stop each particle during the move  $r$ . Each particle has its own row and can be stopped at any column by its obstacle. We leave an empty column between each obstacle to prevent collisions during the next move. **Move 3:** Moving  $d$  arranges the particles into their desired rows. These rows are spread in a staggered horizontal line. **Move 4:** Moving  $l$  stacks the staggered rows into the desired permutation, and returns the array to the initial position.

By reapplying the same permutation enough times, we can return to the original configuration. The permutations shown in Fig. 6 return to the original image in 2 cycles, while Fig. 7 requires 740 cycles. For a two-color image, we can always construct a permutation that resets in 2 cycles. We construct an *involution*, a function that is its own inverse, using cycles of length two that transpose two particles. This technique does not extend to images with more than two colors.



**Fig. 6.** In this image for  $N = 15$ , black cells are obstacles, white cells are free, and colored discs are individual particles. The world has been designed to permute the particles between ‘A’ into ‘B’ every four steps:  $\langle u, r, d, l \rangle$ . See video at <http://youtu.be/3tJdRrNShXM>.



**Fig. 7.** In this larger example with  $N = 100$ , the different control sections are easier to see than in Fig. 6. (1) The staggered obstacles on the left spread the matrix vertically, (2) the scattered obstacles on the upper right permute each element, and (3) the staggered obstacles along the bottom reform each row, which are collected by (4). The cycle resets every 740 iterations. See <http://youtu.be/eExZO0HrWRQ> for an animation of this gadget.

## 5.2 A Workspace for Arbitrary Permutations

There are various ways in which we can exploit Theorem 2 in order to generate larger sets of (or even all) possible permutations. As it turns out, there is a tradeoff between the number of introduced obstacles and the number of moves required for realizing a permutation.


We start with obstacle sets that require only few moves.

**Theorem 3.** *For any set of  $k$  fixed, but arbitrary, permutations of  $n \times n$  pixels, we can construct a set of  $O(kN)$  obstacles, such that we can switch from a start arrangement into any of the  $k$  permutations using at most  $O(\log k)$  force-field moves.*

*Proof.* See Fig. 8. Build a binary tree of depth  $\log k$  for choosing between the permutations by a sequence of  $\langle r, d, (r/l), d, (r/l), \dots, d, (r/l), d, l, u \rangle$  with  $\log k$   $(r/l)$  decisions between the initial prefix  $\langle r, d \rangle$  and final suffix  $\langle d, l, u \rangle$ . This gets the pixels to the set of obstacles for performing the appropriate permutation.

**Corollary 1.** *For any  $\varepsilon > 0$ , we can construct a set of  $(N!)^\varepsilon$  obstacles such that any permutation of  $n \times n = N$  pixels can be achieved by at most  $O(N \log N)$  force-field moves.*

*Proof.* Follows from Theorem 3 by  $k = (N!)^\varepsilon / N$ .



**Fig. 8.** For any set of  $k$  fixed, but arbitrary permutations of  $n \times n$  pixels, we can construct a set of  $O(kN)$  obstacles, such that we can switch from a start arrangement into any of the  $k$  permutations using at most  $O(\log k)$  force-field moves. Here  $k = 4$  and transforms ‘A’ into ‘B’, ‘C’, ‘D’, or ‘E’ in eight moves:  $\langle r, d, (r/l), d, (r/l), d, l, u \rangle$ .

Now we proceed to more economical sets of obstacles, with arbitrary permutations realized by clockwise and counterclockwise move sequences. We make use of the following lemma, which shows that two base permutations are enough to generate any desired rearrangement.

**Lemma 1.** *Any permutation of  $N$  objects can be generated by the two base permutations  $p = (1, 2)$  and  $q = (1, 2, \dots, N)$ . Moreover, any permutation can be generated by a sequence of length at most  $N^2$  that consists of  $p$  and  $q$ .*

*Proof.* Similar to BUBBLE SORT, we use two nested loops of  $N$ . Each move consists of performing  $q$  once, and  $p$  when appropriate.

This allows us to establish the following result.


**Theorem 4.** *We can construct a set of  $O(N)$  obstacles such that any  $n \times n$  arrangement of  $N$  pixels can be rearranged into any other  $n \times n$  arrangement  $\pi$  of the same pixels, using at most  $O(N^2)$  force-field moves.*

*Proof.* See Fig. 9. Use Theorem 2 to build two sets of obstacles, one each for  $p$  and  $q$ , such that  $p$  is realized by the sequence  $\langle u, r, d, l \rangle$  (clockwise) and  $q$  is realized by  $\langle r, u, l, d \rangle$  (counterclockwise). Then we use the appropriate sequence for generating  $\pi$  in  $O(N^2)$  moves.

Using a larger set of generating base permutations allows us to reduce the number of necessary moves. Again, we make use of a simple base set for generating arbitrary permutations.

**Lemma 2.** *Any permutation of  $N$  objects can be generated by the  $N$  base permutations  $p_1 = (1, 2), p_2 = (1, 3), \dots, p_{N-1} = (1, (N-1))$  and  $q = (1, 2 \dots N)$ . Moreover, any permutation can be generated by a sequence of length at most  $N$  that consists of the  $p_i$  and  $q$ .*

*Proof.* Straightforward, analogous to Theorem 4: in each step  $i$ , apply  $q$  once, and swap element  $\pi(i)$  into position  $i$ .



**Fig. 9.** Repeated application of two base permutations can generate any permutation, when used in a manner similar to BUBBLE SORT. The obstacles above generate the base permutation  $p = (1, 2)$  in the clockwise direction  $\langle u, r, d, l \rangle$  and  $q = (1, 2, \dots, N)$  in the counter-clockwise direction  $\langle r, u, l, d \rangle$ .

**Theorem 5.** *We can construct a set of  $O(N^2)$  obstacles such that any  $n \times n$  arrangement of  $N$  pixels can be rearranged into any other  $n \times n$  arrangement  $\pi$  of the same pixels, using at most  $O(N \log N)$  force-field moves.*

*Proof.* Use Theorem 2 to build  $N$  sets of obstacles, one each for  $p_1, \dots, p_{N-1}, q$ . Furthermore, use Lemma 2 for generating all permutations with at most  $N$  different of these base permutation, and Theorem 3 for switching between these  $k = N$  permutations. Then we can get  $\pi$  with at most  $N$  cycles, each consisting of at most  $O(\log N)$  force-field moves.

This is the best possible with respect to the number of moves, in the following sense:

**Theorem 6.** *Suppose we have a set of obstacles such that any permutation of an  $n \times n$  arrangement of pixels can be achieved by at most  $M$  force-field moves. Then  $M$  is at least  $\Omega(N \log N)$ .*

*Proof.* Each permutation must be achieved by a sequence of force-field moves. Because each decision for a force-field move  $\langle u, d, l, r \rangle$  partitions the remaining set of possible permutations into at most four different subsets, we need at least  $\Omega(\log(N!)) = \Omega(N \log N)$  such moves.

## 6 Conclusions

In this paper we analyzed the complexity of steering many particles with uniform inputs in a 2D environment with obstacles. We are motivated by practical challenges in steering magnetically-actuated robots through vascular networks. Many examples of natural, locally 2D vascular networks exist, such as the leaf example in Fig. 1b, and the endothelial networks on the surface of organs.

Clearly, there are many exciting new challenges that lie ahead. The next step is to extend the complexity analysis to PSPACE-COMplete. We are also exploring using particles and obstacles to construct logic gates. These results let us implement AND and OR gates. Using *dual-rail logic*, where the signal and its inverse are explicitly represented by the presence of a particle along

either the signal or the inverse rail, we could also implement NOT, NAND, and NOR gates. Generating fan-out gates seems to require additional complexity in GLOBALCONTROL-MANYPARTICLES because conservation rules are violated. Some way of encoding an order of precedence so that a reversible operation on particle  $a$  will affect particle  $b$  is needed. Potential approaches use either  $2 \times 1$  particles, or  $0.5 \times 1$  obstacles so that the presence of a first particle can enable an action on a second particle, and yet be distinguished from the absence of the first particle and the presence of the second. With uniform  $1 \times 1$  obstacles and particles, these cases are indistinguishable. Finally, many exciting applications require platforms that can navigate in three dimensions. This poses a large number of additional challenges, both for the theory and the physical implementation.

**Acknowledgments.** We acknowledge the helpful discussion and motivating experimental efforts with *T. pyriformis* cells by Yan Ou and Agung Julius at RPI and Paul Kim and MinJun Kim at Drexel University. This work was supported by the National Science Foundation under [CPS-1035716](#).

## References

1. Abelson, H., Allen, D., Coore, D., Hanson, C., Homsy, G., Thomas, J., Knight, F., Nagpal, R., Rauch, E., Sussman, G.J., Weiss, R.: Amorphous computing. *Commun. ACM* **43**(5), 74–82 (2000)
2. Akella, S., Huang, W.H., Lynch, K.M., Mason, M.T.: Parts feeding on a conveyor with a one joint robot. *Algorithmica* **26**(3), 313–344 (2000)
3. Akella, S., Mason, M.T.: Using partial sensor information to orient parts. *Int. J. Robot. Res.* **18**(10), 963–997 (1999)
4. Becker, A., Bretl, T.: Approximate steering of a unicycle under bounded model perturbation using ensemble control. *IEEE Trans. Robot.* **28**(3), 580–591 (2012)
5. Becker, A., Habibi, G., Werfel, J., Rubenstein, M., McLurkin, J.: Massive uniform manipulation: controlling large populations of simple robots with a common input signal. In: *IEEE International Conference on Intelligent Robots and Systems*, pp. 520–527, November 2013
6. Becker, A., Onyuksel, C., Bretl, T.: Feedback control of many differential-drive robots with uniform control inputs. In: *IEEE International Conference on Intelligent Robots and Systems*, October 2012
7. Becker, A., Ou, Y., Kim, P., Kim, M., Julius, A.: Feedback control of many magnetized tetrahymena pyriformis cells by exploiting phase inhomogeneity. In: *IEEE International Conference on Intelligent Robots and Systems*, pp. 3317–3323, November 2013
8. Chanu, A., Felfoul, O., Beaudoin, G., Martel, S.: Adapting the clinical mri software environment for real-time navigation of an endovascular untethered ferromagnetic bead for future endovascular interventions. *Magn. Reson. Med.* **59**(6), 1287–1297 (2008)
9. Chiang, P.-T., Mielke, J., Godoy, J., Guerrero, J.M., Alemany, L.B., Villagómez, C.J., Saywell, A., Grill, L., Tour, J.M.: Toward a light-driven motorized nanocar: Synthesis and initial imaging of single molecules. *ACS Nano* **6**(1), 592–597 (2011)



10. Demaine, E.D., Demaine, M.L., O'Rourke, J.: PushPush and Push-1 are NP-hard in 2D. In: Proceedings of the 12th Annual Canadian Conference on Computational Geometry (CCCG), pp. 211–219, August 2000
11. Demaine, E.D., Hearn, R.A.: Playing games with algorithms: algorithmic combinatorial game theory. In: Albert, M.H., Nowakowski, R.J. (eds.) *Games of No Chance 3*, vol. 56, pp. 3–56. Mathematical Sciences Research Institute Publications, Cambridge University Press, Cambridge (2009)
12. Donald, B.R., Levey, C.G., Paprotny, I., Rus, D.: Planning and control for microassembly of structures composed of stress-engineered mems microrobots. *Int. J. Robot. Res.* **32**(2), 218–246 (2013)
13. Dor, D., Zwick, U.: Sokoban and other motion planning problems. *Comput. Geom.* **13**(4), 215–228 (1999)
14. Engels, B., Kamphans, T.: On the complexity of Randolph's robot game. Technical report, Rheinische Friedrich-Wilhelms-Universität Bonn Institut für Informatik I, University of Cologne, Germany (2005)
15. Erdmann, M., Mason, M.: An exploration of sensorless manipulation. *IEEE J. Robot. Autom.* **4**(4), 369–379 (1988)
16. Fekete S.P., Kröller, A.: Geometry-based reasoning for a large sensor network. In: Proceedings of the 22nd Annual ACM Symposium on Computational Geometry, pp. 475–476 (2006). <http://www.computational-geometry.org/SoCG-videos/socg06video/>
17. Fekete, S.P., Kröller, A.: Topology and routing in sensor networks. In: Kutylowski, M., Cichoń, J., Kubiak, P. (eds.) *ALGOSENSORS 2007*. LNCS, vol. 4837, pp. 6–15. Springer, Heidelberg (2008)
18. Fekete, S.P., Kröller, A., Pfisterer, D., Fischer, S., Buschmann, C.: Neighborhood-based topology recognition in sensor networks. In: Nikolettseas, S.E., Rolim, J. (eds.) *ALGOSENSORS 2004*. LNCS, vol. 3121, pp. 123–136. Springer, Heidelberg (2004)
19. Floyd, S., Diller, E., Pawashe, C., Sitti, M.: Control methodologies for a heterogeneous group of untethered magnetic micro-robots. *Int. J. Robot. Res.* **30**(13), 1553–1565 (2011)
20. Frutiger, D., Kratochvil, B., Vollmers, K., Nelson, B.J.: Magmites - wireless resonant magnetic microrobots. In: IEEE International Conference on Robotics and Automation, Pasadena, CA, May 2008
21. Goemans, O.C., Goldberg, K., van der Stappen, A.F.: Blades: a new class of geometric primitives for feeding 3D parts on vibratory tracks. In: International Conference on Robotics and Automation, pp. 1730–1736, May 2006
22. Goldberg, K., Mirtich, B.V., Zhuang, Y., Craig, J., Carlisle, B.R., Canny, J.: Part pose statistics: estimators and experiments. *IEEE Trans. Robot. Autom.* **15**(5), 849–857 (1999)
23. Goldberg, K.Y.: Orienting polygonal parts without sensors. *Algorithmica* **10**(2), 201–225 (1993)
24. Hearn R.A., Demaine, E.D.: PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. arXiv:cs/0205005, cs.CC/0205005 (2002)
25. Hoffmann, M.: Motion planning amidst movable square blocks: Push-\* is NP-hard. In: Canadian Conference on Computational Geometry, pp. 205–210, June 2000
26. Holzer, M., Schwoon, S.: Assembling molecules in atomix is hard. *Theoret. Comput. Sci.* **313**(3), 447–462 (2004)
27. Kahn, J., Katz, R., Pister, K.: Emerging challenges: mobile networking for smart dust. *J. Comm. Net.* **2**, 188–196 (2000)

28. Khalil, I.S.M., Pichel, M.P., Reefman, B.A., Sukas, O.S., Abelmann, L., Misra, S.: Control of magnetotactic bacterium in a micro-fabricated maze. In: IEEE International Conference on Robotics and Automation, Karlsruhe, Germany, pp. 5488–5493, May 2013
29. Kröller, A., Fekete, S.P., Pfisterer, D., Fischer, S.: Deterministic boundary recognition and topology extraction for large sensor networks. In: Proceedings of the 17th ACM-SIAM Symposium Discrete Algorithms, pp. 1000–1009 (2006)
30. Leighton, F.T.: Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes. Morgan Kaufmann, San Mateo (1991)
31. Lynch, K.M., Northrop, M., Pan, P.: Stable limit sets in a dynamic parts feeder. *IEEE Trans. Robot. Autom.* **18**(4), 608–615 (2002)
32. Moll, M., Erdmann, M.: Manipulation of pose distributions. *Int. J. Robot. Res.* **21**(3), 277–292 (2002)
33. Murphey, T.D., Bernheisel, J., Choi, D., Lynch, K.M.: An example of parts handling and self-assembly using stable limit sets. In: International Conference on Robots and System, pp. 1624–1629, August 2005
34. Murphey, T.D., Burdick, J.W.: Feedback control methods for distributed manipulation systems that involve mechanical contacts. *Int. J. Robot. Res.* **23**(7–8), 763–781 (2004)
35. Ou, Y., Kim, D.H., Kim, P., Kim, M.J., Julius, A.A.: Motion control of magnetized tetrahymena pyriformis cells by magnetic field with model predictive control. *Int. J. Rob. Res.* **32**(1), 129–139 (2013)
36. Peyer, K.E., Zhang, L., Nelson, B.J.: Bio-inspired magnetic swimming microrobots for biomedical applications. *Nanoscale* **5**, 1259–1272 (2013)
37. Rubenstein, M., Ahler, C., Nagpal, R.: Kilobot: a low cost scalable robot system for collective behaviors. In: IEEE International Conference on Robotics and Automation, pp. 3293–3298, May 2012
38. ThinkFun. Tilt: Gravity fed logic maze. <http://www.thinkfun.com/tilt>
39. van der Stappen, A.F., Berretty, R.-P., Goldberg, K., Overmars, M.: Geometry and part feeding. In: Hager, G.D., Christensen, H.I., Bunke, H., Klein, R. (eds.) *Sensor Based Intelligent Robots*. LNCS, vol. 2238, pp. 259–281. Springer, Heidelberg (2002)
40. Vartholomeos, P., Akhavan-Sharif, M., Dupont, P.E.: Motion planning for multiple millimeter-scale magnetic capsules in a fluid environment. In: IEEE International Conference on Robotics and Automation, pp. 1927–1932, May 2012
41. Vose, T.H., Umbanhowar, P., Lynch, K.M.: Friction-induced velocity fields for point parts sliding on a rigid oscillated plate. In: *Robotics: Science and Systems*, Zurich, Switzerland, June 2008
42. Vose, T.H., Umbanhowar, P., Lynch, K.M.: Sliding manipulation of rigid bodies on a controlled 6-DoF plate. *Int. J. Robot. Res.* **31**(7), 819–838 (2012)
43. Wilfong, G.: Motion planning in the presence of movable obstacles. *Ann. Math. Artif. Intell.* **3**(1), 131–150 (1991)