

# 1 TuringMobile: A Turing Machine of Oblivious 2 Mobile Robots with Limited Visibility and its 3 Applications

4 **Giuseppe A. Di Luna**

5 Aix-Marseille University and LiS Laboratory, Marseille, France  
6 giuseppe.diluna@lif.univ-mrs.fr

7 **Paola Flocchini**

8 University of Ottawa, Ottawa, Canada  
9 paola.flocchini@uottawa.ca

10 **Nicola Santoro**

11 Carleton University, Ottawa, Canada  
12 santoro@scs.carleton.ca

13 **Giovanni Viglietta**<sup>1</sup>

14 JAIST, Nomi City, Japan  
15 johnny@jaist.ac.jp

---

## 16 — Abstract —

17 In this paper we investigate the computational power of a set of mobile robots with limited  
18 visibility. At each iteration, a robot takes a snapshot of its surroundings, uses the snapshot to  
19 compute a destination point, and it moves toward its destination. Each robot is punctiform and  
20 memoryless, it operates in  $\mathbb{R}^m$ , it has a local reference system independent of the other robots'  
21 ones, and is activated asynchronously by an adversarial scheduler. Moreover, robots are non-rigid,  
22 in that they may be stopped by the scheduler at each move before reaching their destination (but  
23 are guaranteed to travel at least a fixed unknown distance before being stopped).

24 We show that despite these strong limitations, it is possible to arrange  $3m + 3k$  of these weak  
25 entities in  $\mathbb{R}^m$  to simulate the behavior of a stronger robot that is rigid (i.e., it always reaches  
26 its destination) and is endowed with  $k$  registers of persistent memory, each of which can store  
27 a real number. We call this arrangement a *TuringMobile*. In its simplest form, a TuringMobile  
28 consisting of only three robots can travel in the plane and store and update a single real number.  
29 We also prove that this task is impossible with fewer than three robots.

30 Among the applications of the TuringMobile, we focused on Near-Gathering (all robots have  
31 to gather in a small-enough disk) and Pattern Formation (of which Gathering is a special case)  
32 with limited visibility. Interestingly, our investigation implies that both problems are solvable in  
33 Euclidean spaces of any dimension, even if the visibility graph of the robots is initially discon-  
34 nected, provided that a small amount of these robots are arranged to form a TuringMobile. In  
35 the special case of the plane, a basic TuringMobile of only three robots is sufficient.

36 **2012 ACM Subject Classification** F.2.2 Nonnumerical Algorithms and Problems, I.2.9 Robotics

37 **Keywords and phrases** Mobile Robots, Turing Machine, Real RAM

38 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

39 **Related Version** Full Version on ArXiv: <https://arxiv.org/pdf/1709.08800>

---

<sup>1</sup> [Contact Author. Address: 1-50-D-21 Asahidai, Nomi City, Ishikawa Prefecture 923-1211, Japan. Phone: +81 80-8691-6839.]



## 40 **1** Introduction

### 41 **1.1** Framework and Background.

42 The investigations of systems of autonomous mobile robots have long moved outside the  
 43 boundaries of the engineering, control, and AI communities. Indeed, the computational and  
 44 complexity issues arising in such systems are important research topics within theoretical  
 45 computer science, especially in distributed computing. In these theoretical investigations,  
 46 the robots are usually viewed as computational entities that live in a metric space, typically  
 47  $\mathbb{R}^2$  or  $\mathbb{R}^3$ , in which they can move. Each robot operates in “Look-Compute-Move” (LCM)  
 48 cycles: it observes its surroundings, it computes a destination within the space based on  
 49 what it sees, and it moves toward the destination. The only means of interaction between  
 50 robots are observations and movements: that is, communication is *stigmergic*. The robots,  
 51 identical and outwardly indistinguishable, are *oblivious*: when starting a new cycle, a robot  
 52 has no memory of its activities (observations, computations, and moves) from previous cycles  
 53 (“every time is the first time”).

54 There have been intensive research efforts on the computational issues arising with  
 55 such robots, and an extensive literature has been produced in particular in regard to  
 56 the important class of *Pattern Formation* problems [8, 20, 22, 23, 29, 30] as well as for  
 57 *Gathering* [1, 2, 4, 9, 10, 12, 13, 15, 11, 21, 25] and *Scattering* [6, 24]; see also [7, 14, 31]. The  
 58 goal of the research has been to understand the minimal assumptions needed for a team  
 59 (or swarm) of such robots to solve a given problem, and to identify the impact that specific  
 60 factors have on feasibility and hence computability.

61 The most important factor is the power of the adversarial scheduler that decides when  
 62 each activity of each robot starts and when it ends. The main adversaries (or “environments”)  
 63 considered in the literature are: *synchronous*, in which the computation cycles of all active  
 64 robots are synchronized, and at each cycle either all (in the fully synchronous case) or a  
 65 subset (in the semi-synchronous case) of the robots are activated, and *asynchronous*, where  
 66 computation cycles are not synchronized, each activity can take a different and unpredictable  
 67 amount of time, and each robot can be independently activated at each time instant.

68 An important factor is whether a robot moving toward a computed destination is  
 69 guaranteed to reach it (*rigid* robot), or it can be stopped on the way (*non-rigid*) at a point  
 70 decided by an adversary. In all the above cases, the power of the adversaries is limited by  
 71 some basic fairness assumption. All the existing investigations have concentrated on the  
 72 study of (a-)synchrony, several on the impact of rigidity, some on other relevant factors such  
 73 as agreement on local coordinate systems or on their orientation, etc.; for a review, see [19].

74 From a computational point of view, there is another crucial factor: the visibility range  
 75 of the robots, that is, how much of the surrounding space they are able to observe in a Look  
 76 operation. In this regard, two basic settings are considered: *unlimited visibility*, where the  
 77 robots can see the entire space (and thus all other robots), and *limited visibility*, when the  
 78 robots have a fixed visibility radius. While the investigations on (a-)synchrony and rigidity  
 79 have concentrated on all aspects of those assumptions, this is not the case with respect to  
 80 visibility. In fact, almost all research has assumed unlimited visibility; few exceptions are the  
 81 algorithms for Convergence [4], Gathering [16, 17, 21], and Near-Gathering [25] when the  
 82 visibility range of the robot is limited. The unlimited visibility assumption clearly greatly  
 83 simplifies the computational universe under investigation; at the same time, it neglects the  
 84 more general and realistic one, which is still largely unknown.

85 Let us also stress that, in the existing literature, all results on oblivious robots are for  $\mathbb{R}^1$   
 86 and  $\mathbb{R}^2$ ; the only exception is the recent result on plane formation in  $\mathbb{R}^3$  by semi-synchronous

87 rigid robots with unlimited visibility [31]. No results exist for robots in higher dimensions.

## 88 1.2 Contributions.

89 In this paper we contribute several constructive insights on the computational universe  
90 of oblivious robots with limited visibility, especially asynchronous non-rigid ones, in any  
91 dimension.

### 92 TuringMobile

93 The first and main contribution is the design of a “moving Turing Machine” made solely  
94 of asynchronous oblivious non-rigid robots in  $\mathbb{R}^m$  with limited visibility, for any  $m \geq 2$ .  
95 More precisely, we show how to arrange  $3m + 3k$  identical non-rigid oblivious robots in  $\mathbb{R}^m$   
96 with a visibility radius of  $V + \varepsilon$  (for any  $\varepsilon > 0$ ) and how to program them so that they can  
97 collectively behave as a single rigid robot in  $\mathbb{R}^m$  with  $k$  persistent registers and visibility  
98 radius  $V$  would. This team of identical robots is informally called a *TuringMobile*. We obtain  
99 this result by using as fundamental construction a basic component, which is able to move in  
100  $\mathbb{R}^2$  while storing and updating a single real number. Interestingly, we show that 3 agents  
101 are necessary and sufficient to build such a machine. The TuringMobile will then be built  
102 by arranging multiple copies of this basic component. Notably, the robots that constitute a  
103 TuringMobile need only be able to compute arithmetic operations and square roots.

104 A TuringMobile is a powerful construct that, once deployed in a swarm of robots, can act  
105 as a rigid leader with persistent memory, allowing the swarm to overcome many handicaps  
106 imposed by obliviousness, limited visibility, and asynchrony. As examples we present a  
107 variety of applications in  $\mathbb{R}^m$ , with  $m \geq 2$ .

108 There is a limitation to the use of a TuringMobile when deployed in a swarm of robots.  
109 Namely, the TuringMobile must be always recognizable (e.g., by its unique shape) so that other  
110 robots cannot interfere by moving too close to the machine, disrupting its structure. This  
111 limitation can be overcome when the robots of the TuringMobile are visibly distinguishable  
112 from the others. However, this requirement is not necessary for all applications, but is only  
113 required when we want to perfectly simulate a rigid robot with memory.

114 We remark that we do not discuss how robots can self-assemble into a TuringMobile. We  
115 only focus on how the machine can be designed when we can freely arrange some robots. In  
116 the case of robots with unlimited visibility, a TuringMobile can be self-assembled, provided  
117 that the initial configuration of the robots is asymmetric. In the case of limited visibility,  
118 self-assembling a TuringMobile is a more complex and still open problem.

### 119 Applications

120 We propose several applications of our TuringMobile. First of all, the TuringMobile can  
121 explore and search the space. We then show how it can be employed to solve the long-standing  
122 open problem of (Near-)Gathering with limited visibility in spite of an asynchronous non-  
123 rigid scheduler and disagreement on the axes, a problem still open without a TuringMobile.  
124 Interestingly, the presence of the TuringMobile allows Gathering to be done even if the  
125 initial visibility graph is disconnected. Finally we show how the arbitrary Pattern Formation  
126 problem can be solved under the same conditions (asynchrony, limited visibility, possibly  
127 disconnected visibility graph, etc.).

128 The paper is organized as follows: In Section 2 we give formal definitions, introducing  
129 mobile robots with or without memory as *oracle semi-oblivious real RAMs*. In Section 3  
130 we illustrate our implementation of the TuringMobile. In Section 4 we show how to apply  
131 the TuringMobile to solve fundamental problems. Due to space constraints, the proof of

132 correctness of our TuringMobile implementation, several technical parts of the paper, and  
 133 additional figures can be found in the full paper [18].

## 134 **2** Definitions and Preliminaries

### 135 **2.1 Oracle Semi-Oblivious Real RAMs**

136 **Real random-access machines.** A *real RAM*, as defined by Shamos [26, 28], is a random-  
 137 access machine [3] that can operate on real numbers. That is, instead of just manipulating  
 138 and storing integers, it can handle arbitrary real numbers and do infinite-precision operations  
 139 on them. It has a finite set of internal *registers* and an infinite ordered sequence of *memory*  
 140 *cells*; each register and each memory cell can hold a single real number, which the machine  
 141 can modify by executing its program.<sup>2</sup>

142 A real RAM’s instruction set contains at least the four arithmetic operations, but it may  
 143 also contain  $k$ -th roots, trigonometric functions, exponentials, logarithms, and other analytic  
 144 functions, depending on the application. The machine can also compare two real numbers  
 145 and branch depending on which one is larger.

146 The initial contents of the memory cells are the *input* of the machine (we stipulate that  
 147 only finitely many of them contain non-zero values), and their contents when the machine  
 148 halts are its *output*. So, each program of a real RAM can be viewed as a partial function  
 149 mapping tuples of reals into tuples of reals.

150 **Oracles and semi-obliviousness.** We introduce the *oracle semi-oblivious real RAM*, which  
 151 is a real RAM with an additional “ASK” instruction. Whenever this instruction is executed,  
 152 the contents of all the memory cells are replaced with new values, which are a function of  
 153 the numbers stored in the registers.

154 In other words, the machine can query an external oracle by putting a question in its  $k$   
 155 registers in the form of  $k$  real numbers. The oracle then reads the question and writes the  
 156 answer in the machine’s memory cells, erasing all pre-existing data. The term “semi-oblivious”  
 157 comes from the fact that, every time the machine invokes the oracle, it “forgets” everything  
 158 it knows, except for the contents of the registers, which are preserved.<sup>3</sup>

159 ► **Remark.** In spite of their semi-obliviousness, these real RAMs with oracles are at least as  
 160 powerful as Turing Machines with oracles.

### 161 **2.2 Mobile Robots as Real RAMs**

162 **Mobile robots.** Our oracle semi-oblivious real RAM model can be reinterpreted in the  
 163 realm of *mobile robots*. A mobile robot is a computational entity that lives in a metric space,  
 164 typically  $\mathbb{R}^2$  or  $\mathbb{R}^3$ . It can observe its surroundings and move within the space based on what  
 165 it sees. The same space may be populated by several mobile robots and static objects.

166 To compute its next destination point, a mobile robot executes a real RAM program  
 167 with input a representation of its local view of the space. After moving, its entire memory is  
 168 erased, but the content of its  $k$  registers is preserved. Then it makes a new observation; from

---

<sup>2</sup> Nonetheless, the constant operands in a real RAM’s program cannot be arbitrary real numbers, but have to be integers.

<sup>3</sup> Observe that, in general, the machine cannot salvage its memory by encoding its contents in the registers: since its instruction set has only analytic functions, it cannot injectively map a tuple of arbitrary real numbers into a single real number.

169 the observation data and the contents of the registers, it computes another destination point,  
170 and so on. If  $k = 0$ , the mobile robot is said to be *oblivious*.

171 The actual movement of a mobile robot is controlled by an external *scheduler*. The  
172 scheduler decides how fast the robot moves toward its destination point, and it may even  
173 interrupt its movement before the destination point is reached. If the movement is interrupted  
174 midway, the robot makes the next observation from there and computes a new destination  
175 point as usual. The robot is not notified that an interruption has occurred, but it may be  
176 able to infer it from its next observation and the contents of its registers. For fairness, the  
177 scheduler is only allowed to interrupt a robot after it has covered a distance of at least  $\delta$  in  
178 the current movement, where  $\delta$  is a positive constant. This guarantees, for example, that if  
179 a robot keeps computing the same destination point, it will reach it in a finite number of  
180 iterations. If  $\delta = \infty$ , the robot always reaches its destination, and is said to be *rigid*.

181 **Mobile robots, revisited.** A mobile robot in  $\mathbb{R}^m$  with  $k$  registers can be modeled as an  
182 oracle semi-oblivious real RAM with  $2m + k + 1$  registers, as follows.

- 183 ■  $m$  *position registers* hold the absolute coordinates of the robot in  $\mathbb{R}^m$ .
- 184 ■  $m$  *destination registers* hold the destination point of the robot, expressed in its local  
185 coordinate system.
- 186 ■ 1 *timestamp register* contains the time of the robot's last observation.
- 187 ■  $k$  *true registers* correspond to the registers of the robot.

188 As the RAM's execution starts, it ignores its input, erases all its registers, and executes  
189 an "ASK" instruction. The oracle then fills the RAM's memory with the robot's initial  
190 position  $p$ , the time  $t$  of its first observation, and a representation of the geometric entities  
191 and objects surrounding the robot, as seen from  $p$  at time  $t$ .

192 The RAM first copies  $p$  and  $t$  in its position registers and timestamp register, respectively.  
193 Then it executes the program of the mobile robot, using its true registers as the robot's  
194 registers and adding  $m + 1$  to all memory addresses. This effectively makes the RAM ignore  
195 the values of  $p$  and  $t$ , which indeed are not supposed to be known to the mobile robot.

196 When the robot's program terminates, the RAM's memory contains the output, which is  
197 the next destination point  $p'$ , expressed in the robot's coordinate system. The RAM copies  $p'$   
198 into its destination registers, and the execution jumps back to the initial "ASK" instruction.

199 Now the oracle reads  $p$ ,  $p'$ , and  $t$  from the RAM's registers (it ignores the true registers),  
200 converts  $p'$  in absolute coordinates (knowing  $p$  and the orientation of the local coordinate  
201 system of the robot) and replies with a new position  $p''$ , a timestamp  $t' > t$ , and observation  
202 data representing a snapshot taken from  $p''$  at time  $t'$ . To comply with the mobile robot  
203 model,  $p''$  must be on the segment  $pp'$ , such that either  $p'' = p'$  or  $\overline{pp''} \geq \delta$ . The execution  
204 then proceeds in the same fashion, indefinitely.

205 Note that in this setting the oracle represents the scheduler. The presence of a timestamp  
206 in the query allows the oracle to model dynamic environments in which several independent  
207 robots may be moving concurrently (without a timestamp, two observations from the same  
208 point of view would always be identical).

209 **Snapshots and limited visibility.** In the mobile robot model we consider in this paper,  
210 an observation is simply an instantaneous *snapshot* of the environment taken from the robot's  
211 position. In turn, each entity and object that the robot can see is modeled as a dimensionless  
212 point in  $\mathbb{R}^m$ . A mobile robot has a positive *visibility radius*  $V$ : it can see a point in  $\mathbb{R}^m$  if  
213 and only if it is located at distance at most  $V$  from its current position. If  $V = \infty$ , the robot  
214 is said to have *unlimited visibility*.

215 As we hinted at earlier in this section, a mobile robot has its own local reference system  
 216 in which all the coordinates of the objects in its snapshots are expressed. The origin of a  
 217 robot's local coordinate system always coincides with the robot's position (hence it follows  
 218 the robot as it moves), and its orientation and handedness are decided by the scheduler  
 219 (and remain fixed). Different mobile robots may have coordinate systems with a different  
 220 orientation or handedness. (However, when two robots have the same visibility radius, they  
 221 also implicitly have the same unit of distance.)

222 So, a snapshot is just a (finite) list of points, each of which is an  $m$ -tuple of real numbers.

223 **Simulating memory and rigidity.** The main contribution of this paper, loosely speaking,  
 224 is a technique to turn non-rigid oblivious robots into rigid robots with persistent memory,  
 225 under certain conditions. More precisely, if  $3m + 3k$  identical non-rigid oblivious robots in  
 226  $\mathbb{R}^m$  with a visibility radius of  $V + \varepsilon$  (for any  $\varepsilon > 0$ ) are arranged in a specific pattern and  
 227 execute a specific algorithm, they can collectively act in the same way as a single rigid robot  
 228 in  $\mathbb{R}^m$  with  $k > 0$  persistent registers and visibility radius  $V$  would. This team of identical  
 229 robots is informally called a *TuringMobile*.


230 We stress that the robots of a TuringMobile are *asynchronous*, that is, the scheduler  
 231 makes them move at independent arbitrary speeds, and each robot takes the next snapshot  
 232 an arbitrary amount of time after terminating each move. The robots are also *anonymous*,  
 233 in that they are indistinguishable from each other, and they all execute the same program.

234 Although our technique is fairly general and has a plethora of concrete applications  
 235 (some are discussed in Section 4), a “perfect simulation” is achieved only under additional  
 236 conditions on the scheduler or on the environment (see Section 3.2).

### 237 **3 Implementing the TuringMobile**

#### 238 **3.1 Basic Implementation**

239 We will first describe how to construct a basic version of the TuringMobile with just three  
 240 oblivious non-rigid robots in  $\mathbb{R}^2$ . This TuringMobile can remember a single real number  
 241 and rigidly move in the plane by fixed-length steps: its layout is sketched in Figure 1. In  
 242 Section 3.2, we will show how to combine several copies of this basic machine to obtain a  
 243 full-fledged TuringMobile.



■ **Figure 1** Basic TuringMobile at rest, not drawn to scale ( $\mu$  and  $\lambda$  should be smaller)

244 **Position at rest.** The elements of the basic TuringMobile are three: a *Commander* robot,  
 245 a *Number* robot, and a *Reference* robot, located in  $C$ ,  $N$ , and  $R$ , respectively. These robots  
 246 have the same visibility radius of  $V + \varepsilon$ , where  $\varepsilon \ll V$ , and there is always a disk of radius  $\varepsilon$   
 247 containing all three of them. When the machine is “at rest”,  $\angle NRC$  is a right angle, the  
 248 distance between  $C$  and  $R$  is some fixed value  $d \ll \varepsilon$ , and the distance between  $R$  and  $N$  is  
 249 approximately  $2d$ . More precisely,  $N$  lies on a segment  $QQ'$  of length  $\lambda$ , where  $\lambda \ll d$  is some  
 250 fixed value, such that  $Q$  has distance  $2d - \lambda/2$  from  $R$  and  $Q'$  has distance  $2d + \lambda/2$  from  $R$ .

251 **Representing numbers.** The distance between the Reference robot and the Number  
 252 robot when the TuringMobile is at rest is a representation of the real number  $r$  that  
 253 the machine is currently storing. One possible technique is to encode the number  $r$  as  
 254  $\overline{RN} = 2d + \arctan(r) \cdot \lambda/\pi$  and to decode it as  $r = \tan((\overline{RN} - 2d) \cdot \pi/\lambda)$ . However, there  
 255 are also more complicated methods that use only arithmetic functions (see the full paper  
 256 [18]).

257 **Movement directions.** The Commander’s role is to decide in which direction the machine  
 258 should move next, and to initiate the movement. When the machine is at rest, the Commander  
 259 may choose among three possible *final destinations*, labeled  $D_1$ ,  $D_2$ , and  $D_3$  in Figure 1.  
 260 The segments  $CD_1$ ,  $CD_2$ , and  $CD_3$  all have the same length  $\mu$ , with  $\lambda \ll \mu \ll d$ , and form  
 261 angles of  $120^\circ$  with one another, in such a way that  $D_1$  is collinear with  $R$  and  $C$ .

262 Around the center of each segment  $CD_i$  there is a *midway triangle*  $\tau_i$ , drawn in gray in  
 263 Figure 1. This is an isosceles triangle of height  $\lambda$  whose base lies on  $CD_i$  and has length  $\lambda$   
 264 as well. When the Commander decides that its final destination is  $D_i$ , it moves along the  
 265 segment  $CD_i$ , but it takes a detour in the midway triangle  $\tau_i$ , as we will explained shortly.

266 **Structure of the algorithm.** Algorithm 1 is the program that each element of the basic  
 267 TuringMobile executes every time it computes its next destination point.

---

#### Algorithm 1 Basic TuringMobile in $\mathbb{R}^2$

---

```

1: Identify Commander, Number, Reference (located in  $C$ ,  $N$ ,  $R$ , respectively)
2: if I am Commander then
3:   Compute Virtual Commander  $C'$  (based on  $R$  and  $N$ ) and points  $A_i, S_i, S'_i, B_i, D_i$ 
4:   if I am in  $C'$  then Choose final destination  $D_i$  and move to  $A_i$ 
5:   else if  $\exists i \in \{1, 2, 3\}$  s.t. I am on segment  $C'A_i$  but not in  $A_i$  then Move to  $A_i$ 
6:   else if  $\exists i \in \{1, 2, 3\}$  s.t. I am in  $A_i$  then
7:     Move to point  $P$  on segment  $S_iS'_i$  such that  $\overline{PS_i} = f(\overline{NQ})$ 
8:   else if  $\exists i \in \{1, 2, 3\}$  s.t. I am in triangle  $A_iS_iS'_i$  but not on segment  $S_iS'_i$  then
9:     Move to the intersection of segment  $S_iS'_i$  with the extension of line  $A_iC$ 
10:  else if  $\exists i \in \{1, 2, 3\}$  s.t. I am on  $S_iS'_i$  and  $\overline{NQ} = \overline{CS_i}$  then Move to  $B_i$ 
11:  else if  $\exists i \in \{1, 2, 3\}$  s.t. I am in triangle  $B_iS_iS'_i$  but not in  $B_i$  then Move to  $B_i$ 
12:  else if  $\exists i \in \{1, 2, 3\}$  s.t. I am on segment  $B_iD_i$  but not in  $D_i$  then Move to  $D_i$ 
13: else if I am Number then
14:   if  $\overline{CR} = d + \mu$  or  $\overline{CR} = d'$  then
15:     Compute Virtual Commander  $C'$  (based on  $C$  and  $R$ ) and points  $D'_i$ 
16:     if  $\overline{CR} = d + \mu$  and I am not in  $D'_1$  then Move to  $D'_1$ 
17:     else if  $\overline{CR} = d'$  and  $\angle NRC > 90^\circ$  and I am not in  $D'_2$  then Move to  $D'_2$ 
18:     else if  $\overline{CR} = d'$  and  $\angle NRC < 90^\circ$  and I am not in  $D'_3$  then Move to  $D'_3$ 
19:   else
20:     Compute Virtual Commander  $C'$  (based on  $R$  and  $N$ ) and points  $S_i, S'_i$ 
21:     if  $\exists i \in \{1, 2, 3\}$  s.t.  $C$  is on segment  $S_iS'_i$  then
22:       Move to point  $P$  on segment  $QQ'$  such that  $\overline{PQ} = \overline{CS_i}$ 
23:   else if I am Reference then
24:     if Commander and Number are not tasked to move (based on the above rules) then
25:        $\gamma$  = circle centered in  $C$  with radius  $d$ 
26:        $\gamma'$  = circle with diameter  $CN$ 
27:       Move to the intersection of  $\gamma$  and  $\gamma'$  closest to  $R$ 

```

---

268 Since the robots are anonymous, they first have to determine their roles, i.e., who is the  
 269 Commander, etc. (line 1 of the algorithm). We make the assumption that there exists a disk

270 of radius  $\varepsilon$  containing only the TuringMobile (close to its center) and no other robot. Using  
 271 the fact that the two closest robots must be the Commander and the Reference robot and  
 272 that the two farthest robots must be the Commander and the Number robot, it is then easy  
 273 to determine who is who (these properties will be preserved throughout the execution, as  
 274 proved in the full paper [18]).

275 Once it has determined its role, each robot executes a different branch of the algorithm  
 276 (cf. lines 2, 13, and 23). The general idea is that, when the Commander realizes that the  
 277 machine is in its rest position, it decides where to move next, i.e., it chooses a final destination  
 278  $D_i$ . This choice is based on the number  $r$  stored in the machine's "memory" (i.e., the number  
 279 encoded by  $\overline{RN}$ ), the relative positions of the visible robots external to the machine, and  
 280 also on the application, i.e., the specific program that the TuringMobile is executing.

281 When the Commander has decided its final destination  $D_i$ , the entire machine moves by  
 282 the vector  $\overrightarrow{CD_i}$ , and the Number robot also updates its distance from the Reference robot to  
 283 represent a different real number  $r'$ . Again, this number is computed based on the number  $r$   
 284 the machine was previously representing, the relative positions of the visible robots external  
 285 to the machine, and the specific program: in general, the new distance between  $N$  and  $Q$  is  
 286 a function  $f$  of the old distance. When all this is done, the machine is in its rest position  
 287 again, so the Commander chooses a new destination, and so on, indefinitely.

288 **Coordinating movements.** Note that it is not possible for all three robots to translate by  
 289  $\overrightarrow{CD_i}$  at the same time, because they are non-rigid and asynchronous. If the scheduler stops  
 290 them at arbitrary points during their movement, after the structure of the machine has been  
 291 destroyed, they will be incapable of recovering all the information they need to resume their  
 292 movement (recall that they are oblivious and they have to compute a destination point from  
 293 scratch every time).


294 To prevent this, the robots employ various coordination techniques. First the Commander  
 295 moves to the middle triangle  $\tau_i$ , and precisely to its base vertex  $A_i$ , as shown in Figure 2(a)  
 296 (cf. line 5 of Algorithm 1). Then it positions itself on the altitude  $S_iS'_i$ , in such a way as  
 297 to indicate the new number  $r'$  that the machine is supposed to represent. That is, the  
 298 Commander picks the point on  $S_iS'_i$  at distance  $f(\overline{NQ})$  from  $S_i$  (lines 6 and 7). Even if it  
 299 is stopped by the scheduler before reaching such a point, it can recover its destination by  
 300 drawing a ray from  $A_i$  to its current position and intersecting it with  $S_iS'_i$  (lines 8 and 9).

301 When the Commander has reached  $S_iS'_i$ , it waits to let the Number robot adjust its  
 302 position on the segment  $QQ'$  to match that of the Commander on  $S_iS'_i$ , as in Figure 2(b)  
 303 (lines 21 and 22). This effectively makes the Number robot represent the new number  $r'$ .  
 304 Note that the Number robot can do this even if it is stopped by the scheduler several times  
 305 during its march, because the Commander keeps reminding it of the correct value of  $r'$ : since  
 306  $r'$  depends on the old number  $r$ , the Number robot would be unable to re-compute  $r'$  after  
 307 it has forgotten  $r$ . Once the Number robot has reached the correct position on  $QQ'$ , the  
 308 Commander starts moving again (line 10) and finally reaches  $D_i$  while the other robots wait,  
 309 as in Figure 2(c) (lines 11 and 12).

310 When the Commander has reached  $D_i$ , the Number robot realizes it and makes the  
 311 corresponding move (lines 14–18) while the other two robots wait. The destination point of  
 312 the Number robot is  $D'_i$ , as shown in Figure 1. Finally, when the Number robot is in  $D'_i$ , the  
 313 Reference robot realizes it and makes the final move to bring the TuringMobile back into a  
 314 rest position (lines 23–27).

315 **Computing the Virtual Commander.** After the Commander has left its rest position  
 316 and is on its way to  $D_i$ , the TuringMobile loses its initial shape, and identifying the  $D_i$ 's and





■ **Figure 2** Coordinated movement of the Commander and the Number robot

317 the midway triangles becomes non-trivial. So, the robots try to guess where the Commander's  
 318 original rest position may have been by computing a point called the *Virtual Commander*  $C'$ .

319 Assuming that the Reference and Number robots have not moved from their rest positions,  
 320 the Virtual Commander is easily computed: draw the line  $\ell$  through  $R$  perpendicular to  $RN$ ;  
 321 then,  $C'$  is the point on  $\ell$  at distance  $d$  from  $R$  that is closest to  $C$ . Once we have  $C'$ , we  
 322 can construct the points  $D_i$  with respect to  $C'$  (in the same way as we did in Figure 1 with  
 323 respect to  $C$ ). This technique is used by Algorithm 1 at lines 3 and 20.

324 In the special case where the Commander has reached its final destination  $D_i$  and the  
 325 Reference robot has not moved from its rest position (but perhaps the Number robot has  
 326 moved), the Virtual Commander can also be computed. This situation is recognized because  
 327 the distance between the Commander and the Reference robot is either maximum (i.e.,  $d + \mu$ )  
 328 or minimum (i.e.,  $d' = \sqrt{d^2 + \mu^2 - d\mu}$ ), as Figure 1 shows. If the distance is maximum, then  
 329  $C$  must coincide with  $D_1$ ; otherwise,  $C$  coincides with  $D_2$  (if the angle  $\angle NRC$  is obtuse) or  
 330  $D_3$  (if the angle  $\angle NRC$  is acute). Since we know the position of  $R$  and one of the  $D_i$ 's, it is  
 331 then easy to determine the other  $D_i$ 's. This technique is used at line 15.

332 **The Reference robot's behavior.** To know when it has to start moving, the Reference  
 333 robot executes Algorithm 1 from the perspective of the Commander and the Number robot:  
 334 if neither of them is supposed to move, then the Reference robot starts moving (line 24).

335 We have seen that the Number robot can determine its destination  $D'_i$  solely by looking  
 336 at the positions of  $C$  and  $R$ , which remain fixed as it moves. For the Reference robot the  
 337 destination point is not as easy to determine, because the distance between  $C$  and  $N$  varies  
 338 depending on what number is stored in the TuringMobile.

339 However, the Reference robot knows that its move must put the TuringMobile in a rest  
 340 position. The condition for this to happen is that its destination point be at distance  $d$  from  
 341  $C$  (line 25) and form a right angle with  $C$  and  $N$  (line 26). There are exactly two such  
 342 points in the plane, but one of them has distance much greater than  $\mu$  from  $R$ , and hence  
 343 the Reference robot will pick the other (line 27).

344 As the Reference robot moves toward such a point, all the above conditions must be  
 345 preserved, due to the asynchronous and non-rigid nature of the robots. This is not a trivial  
 346 requirement, and a proof that it is indeed fulfilled is in the full paper [18].

### 347 3.2 Complete Implementation

348 We have shown how to implement a basic component of the TuringMobile in  $\mathbb{R}^2$  consisting  
 349 of three robots: a Commander, a Number, and a Reference. The basic component is able to  
 350 rigidly move by a fixed distance  $\mu$  in three fixed directions,  $120^\circ$  apart from one another. It  
 351 can also store and update a single real number.

352 **Planar layout.** We can obtain a full-fledged TuringMobile in  $\mathbb{R}^2$  by putting several tiny  
 353 copies of the basic component side by side. For the machine to work, we stipulate that there  
 354 exists a disk of radius  $\sigma$  that contains all the robots constituting the TuringMobile and no  
 355 extraneous robot, where  $\sigma \ll \varepsilon$ . The distance between two consecutive basic components of  
 356 the TuringMobile is roughly  $s$ , where  $d \ll s \ll \sigma$ . This makes it easy for the robots to tell  
 357 the basic components apart and determine the role of each robot within its basic component.

358 Since a basic component of the TuringMobile is a scalene triangle, which is chiral, all its  
 359 members implicitly agree on a clockwise direction even if they have different handedness.  
 360 Similarly, all robots in the Turing Mobile agree on a “leftmost” basic component, whose  
 361 Commander is said to be the *Leader* of the whole machine.

362 **Coordinated movements.** All the basic components of the TuringMobile are always  
 363 supposed to agree on their next move and proceed in a roughly synchronous way. To achieve  
 364 this, when all the basic components are in a rest position, the Leader decides the next  
 365 direction among the three possible, and executes line 4 of Algorithm 1. Then all the other  
 366 Commanders see where the Leader is going, and copy its movement.

367 When all the Commanders are in their respective  $A_i$ 's, they execute line 7 of the algorithm,  
 368 and so on. At any time, each robot executes a line of the algorithm only if all its homologous  
 369 robots in the other basic components of the TuringMobile are ready to execute that line or  
 370 have already executed it; otherwise, it waits. When the last Reference robot has completed  
 371 its movement, the machine is in a rest position again, and the coordinated execution repeats  
 372 with the Leader choosing another direction, etc.

373 **Simulating a non-oblivious rigid robot.** Let a program for a rigid robot  $\mathcal{R}$  in  $\mathbb{R}^2$  with  
 374  $k$  persistent registers and visibility radius  $V$  be given. We want the TuringMobile described  
 375 above to act as  $\mathcal{R}$ , even though its constituting robots are non-rigid and oblivious.

376 Our TuringMobile consists of  $2 + k$  basic components, each dedicated to memorizing and  
 377 updating one real number. These  $2 + k$  numbers are the  $x$  coordinate and the  $y$  coordinate  
 378 of the destination point of  $\mathcal{R}$  and the contents of the  $k$  registers of  $\mathcal{R}$ . We will call the first  
 379 two numbers the  $x$  variable and the  $y$  variable, respectively.

380 When the TuringMobile is in a rest position, its  $x$  and  $y$  variables represent the co-  
 381 ordinates of the destination point of  $\mathcal{R}$  relative to the Leader of the machine. Whenever  
 382 the TuringMobile moves by  $\mu$  in some direction, these values are updated by subtracting  
 383 the components of an appropriate vector of length  $\mu$  from them. Of course, this update  
 384 is computed by the Commanders of the first two basic components of the machine, which  
 385 communicate it to their respective Number robots, as explained in Section 3.1.

386 Let  $P$  be the destination point of  $\mathcal{R}$ . Since the TuringMobile can only move by vectors of  
 387 length  $\mu$  in three possible directions, it may be unable to reach  $P$  exactly. So, the Leader  
 388 always plans the next move trying to reduce its distance from  $P$  until this distance is at  
 389 most  $2\sigma$  (this is possible because  $\mu \ll d \ll \sigma$ ).

390 When the Leader is close enough to  $P$ , it “pretends” to be in  $P$ , and the TuringMobile  
 391 executes the program of  $\mathcal{R}$  to compute the next destination point. Recall that the visibility  
 392 radius of  $\mathcal{R}$  is  $V$ , and that of the robots of the TuringMobile is  $V + \varepsilon$ . Since  $\sigma \ll \varepsilon$ , each  
 393 member of the TuringMobile can therefore see everything that would be visible to  $\mathcal{R}$  if it

394 were in  $P$ , and compute the output of the program of  $\mathcal{R}$  independently of the other members.  
 395 The only thing it should do when it executes the program of  $\mathcal{R}$  is subtract the values of the  
 396  $x$  and  $y$  variables to everything it sees in its snapshot, discard whatever has distance greater  
 397 than  $V$  from the center, and of course discard the robots of the TuringMobile and replace  
 398 them with a single robot in the center. Then, the robots that are responsible for updating  
 399 the  $x$  and  $y$  variables add the relative coordinates of the new destination point of  $\mathcal{R}$  to these  
 400 variables. Similarly, the robots responsible for updating the  $k$  registers of  $\mathcal{R}$  do so.

401 **Restrictions.** The above TuringMobile correctly simulates  $\mathcal{R}$  under certain conditions.  
 402 The first one is that, if all robots are indistinguishable, then no robot extraneous to the  
 403 TuringMobile may get too close to it (say, within a distance of  $\sigma$  of any of its members). This  
 404 kind of restriction cannot be dispensed with: whatever strategy a team of oblivious robots  
 405 employs to simulate a single non-oblivious robot's behavior is bound to fail if extraneous  
 406 robots join the team creating ambiguities between its members. Nevertheless, the restriction  
 407 can be removed if the members of a TuringMobile are distinguishable from all other robots.

408 Another difficulty comes from the fact that, if the TuringMobile is made of more than one  
 409 basic component and its Commanders are all in their respective  $A_i$ 's and ready to update  
 410 the values represented by the machine, they may get their screenshots at different times,  
 411 due to asynchrony. If the environment moves in the meantime, the screenshots they get are  
 412 different, and this may cause the machine to compute an incorrect destination point or put  
 413 inconsistent values in its simulated registers.

414 There are several possible solutions to this problem: we will only mention two trivial  
 415 ones. We could assume the Commanders to be *synchronous*, that is, make the scheduler  
 416 activate them in such a way that all of them take their screenshots at the same time. This  
 417 way, all Commanders get compatible screenshots and compute consistent outputs. Another  
 418 possible solution is to make the TuringMobile operate in an environment where everything  
 419 else is static, i.e., no moving entities are present other than the TuringMobile's members.

420 We stress that these restrictions make sense if a perfect simulation of  $\mathcal{R}$  is sought. As  
 421 we will see in Section 4, there are several other applications of the TuringMobile technique  
 422 where no such restriction is required.

423 **Higher dimensions.** Let us now generalize the above construction of a planar TuringMobile  
 424 to  $\mathbb{R}^m$ , for any  $m \geq 2$ . We start with the same TuringMobile  $\mathcal{M}$  with  $2 + k$  basic components  
 425 laid out on a plane  $\gamma \subset \mathbb{R}^m$ . Since  $\mathcal{M}$  has only two basic components for the  $x$  and  $y$   
 426 variables, we will add  $m - 2$  basic components to it, positioned as follows.

427 Let vectors  $v_1$  and  $v_2$  be two orthonormal generators of  $\gamma$ , and let us complete  $\{v_1, v_2\}$  to  
 428 an orthonormal basis  $\{v_1, v_2, \dots, v_m\}$  of  $\mathbb{R}^m$ . Now, for all  $i \in \{3, 4, \dots, m\}$ , we make a copy  
 429 of the basic component of  $\mathcal{M}$  containing the Leader, we translate it by  $s \cdot v_i$ , and we add it  
 430 to the TuringMobile ( $s$  is the same value used in the construction of the planar TuringMobile  
 431 at the beginning of Section 3.2). Note that the Leader of this new TuringMobile  $\mathcal{M}'$  is still  
 432 easy to identify, as well as the plane  $\gamma$  when  $\mathcal{M}'$  is at rest.

433 Clearly,  $m$  basic components allow the machine to record a destination point in  $\mathbb{R}^m$ , as  
 434 opposed to  $\mathbb{R}^2$ . Additionally, the positions of the basic components with respect to  $\gamma$  give  
 435 the machine an  $m$ -dimensional sense of direction (see the full paper [18] for further details).

436 ► **Theorem 1.** *Under the aforementioned restrictions, a rigid robot in  $\mathbb{R}^m$  with  $k$  persistent*  
 437 *registers and visibility radius  $V$  can be simulated by a team of  $3m + 3k$  non-rigid oblivious*  
 438 *robots in  $\mathbb{R}^m$  with visibility radius  $V + \varepsilon$ .* ◀

439 **4 Applications**

440 In this section we discuss some applications of the TuringMobile. We also prove that the  
 441 basic TuringMobile constructed in Section 3.1 is minimal, in the sense that no smaller team  
 442 of oblivious robots can accomplish the same tasks.


443 **4.1 Exploring the Plane**

444 The first elementary task a basic TuringMobile in  $\mathbb{R}^2$  can fulfill is that of *exploring* the  
 445 plane. The task consists in making all the robots in the TuringMobile see every point in  
 446 the plane in the course of an infinite execution. We first assume that the three members of  
 447 the TuringMobile are the only robots in the plane. Later in this section, we will extend our  
 448 technique to other types of scenarios and more complex tasks.

449 ► **Theorem 2.** *A basic TuringMobile consisting of three robots in  $\mathbb{R}^2$  can explore the plane.*

450 **Proof.** Recall that a basic TuringMobile can store a single real number  $r$  and update it at  
 451 every move as a result of executing a real RAM program with input  $r$ . In particular, the  
 452 TuringMobile can count how many times it has moved by simply starting its execution with  
 453  $r = 0$  and computing  $r := r + 1$  at each move.

454 Moreover, the Commander chooses the direction of the next move (in the form of a point  
 455  $D_i$ , see Figure 1) by executing another real RAM program with input  $r$ . If  $r$  is an integer,  
 456 the Commander can therefore compute any Turing-computable function on  $r$ , and so it can  
 457 decide to move to  $D_1$  the first time, then to  $D_2$  twice, then to  $D_3$  three times, to  $D_1$  four  
 458 times, and so on. This pattern of moves is illustrated in Figure 3, and of course it results in  
 459 the exploration of the plane, because the visibility radius of the robots is much greater than  
 460 the step  $\mu$ . ◀



■ **Figure 3** Exploration of the plane by a basic TuringMobile

461 **4.2 Minimality of the Basic TuringMobile**

462 We can use the previous result to prove indirectly that our basic TuringMobile design is  
 463 minimal, because no team of fewer than three oblivious robots in  $\mathbb{R}^2$  can explore the plane.

464 ► **Theorem 3.** *If only one or two oblivious identical robots with limited visibility are present*  
 465 *in  $\mathbb{R}^2$ , they cannot explore the plane, even if the scheduler lets them move synchronously and*  
 466 *rigidly.*

467 **Proof.** Assume that a single oblivious robot is given in  $\mathbb{R}^2$ . Since it always gets the same  
 468 snapshot, it always computes the same destination point in its local coordinate system, and  
 469 so it always translates by the same vector. As a consequence, it just moves along a straight  
 470 ray, and therefore it cannot explore the plane.

471 Let two oblivious robots be given, and suppose that their local coordinate systems are  
 472 oriented symmetrically. Whether the robots see each other or not, if they take their snapshots  
 473 simultaneously, they get identical views, and so they compute destination points that are  
 474 symmetric with respect to  $O$ . If they keep moving synchronously and rigidly,  $O$  remains  
 475 their midpoint. So, if the robots have visibility radius  $V$ , they see each other if and only if  
 476 they are in the circle  $\gamma$  of radius  $V/2$  centered in  $O$ .

477 Let  $O$  be the midpoint of the robots' locations, and consider a Cartesian coordinate  
 478 system with origin  $O$ . Without loss of generality, when the robots do not see each other,  
 479 they move by vectors  $(1, 0)$  and  $(-1, 0)$ , respectively. Let  $\xi$  be the half-plane  $y \geq V$ , and  
 480 observe that  $\xi$  lies completely outside  $\gamma$ .

481 It is obvious that the robots cannot explore the entire plane if neither of them ever stops  
 482 in  $\xi$ . The first time one of them stops in  $\xi$ , it takes a snapshot from there, and starts moving  
 483 parallel to the  $x$  axis, thus never seeing the other robot again, and never leaving  $\xi$ . Of course,  
 484 following a straight line through  $\xi$  is not enough to explore all of it. ◀

### 485 4.3 Near-Gathering with Limited Visibility

486 The exploration technique can be applied to several more complex problems. The first we  
 487 describe is the *Near-Gathering* problem, in which all robots in the plane must get in the  
 488 same disk of a given radius  $\varepsilon$  (without colliding) and remain there forever. It does not matter  
 489 if the robots keep moving, as long as there is a disk of radius  $\varepsilon$  that contains them all.

490 It is clear that solving this problem from every initial configuration is not possible, and  
 491 hence some restrictive assumptions have to be made. The usual assumption is that the initial  
 492 visibility graph of the robots be connected [21, 25]. Here we make a different assumption:  
 493 there are three robots that form a basic TuringMobile somewhere in the plane, and each robot  
 494 not in the TuringMobile has distance at least  $\varepsilon$  from all other robots. (Actually we could  
 495 weaken this assumption much more, but this simple example is good enough to showcase our  
 496 technique.)

497 Say that all robots in the plane have a visibility radius of  $V \gg \varepsilon$ , and that the TuringMobile  
 498 moves by  $\mu \ll \varepsilon$  at each step. The TuringMobile starts exploring the plane as above, and  
 499 it stops in a rest position as soon as it finds a robot whose distance from the Commander  
 500 is smaller than  $V/2$  and greater than  $\varepsilon$ . On the other hand, if a robot is not part of the  
 501 TuringMobile, it waits until it sees a TuringMobile in a rest position at distance smaller than  
 502  $V/2$ . When it does, it moves to a designated area  $\mathcal{A}$  in the proximity of the Commander.  
 503 Such an area has distance at least  $3d$  from the Commander, so no confusion can arise in  
 504 the identification of the members of the TuringMobile. If several robots are eligible to move  
 505 to  $\mathcal{A}$ , only one at a time does so: note that the layout of the TuringMobile itself gives an  
 506 implicit total order to the robots around it. Observe that the robots cannot form a second  
 507 TuringMobile while they move to  $\mathcal{A}$ : in order to do so, two of them would have to move to  
 508  $\mathcal{A}$  at the same time and get close enough to a third robot. Once they enter  $\mathcal{A}$ , the robots  
 509 position themselves on a segment much shorter than  $d$ , so they cannot possibly be mistaken

510 for a TuringMobile.

511 Once the eligible robots have positioned themselves in  $\mathcal{A}$ , the TuringMobile resumes its  
 512 exploration of the plane, and the robots in  $\mathcal{A}$  copy all its movements. Now, if the total  
 513 number of robots in the plane is known, the TuringMobile can stop as soon as all of them  
 514 have joined it. Otherwise, the machine simply keeps exploring the plane forever, eventually  
 515 collecting all robots. In both cases, the Near-Gathering problem is solved.

#### 516 4.4 Pattern Formation with Limited Visibility

517 Suppose the robots are exactly  $n$ , and they are tasked to form a given *pattern* consisting of a  
 518 multiset of  $n$  points: this is the *Pattern Formation* problem, which becomes the *Gathering*  
 519 problem in the special case in which the points are all coincident. For this problem, it does  
 520 not matter where the pattern is formed, nor does its orientation or scale.

521 Again, the Pattern Formation problem is unsolvable from some initial configurations, so  
 522 we make the same assumptions as with the Near-Gathering problem. The algorithm starts  
 523 by solving the Near-Gathering problem as before. The only difference is that now there is a  
 524 second tiny area  $\mathcal{B}$ , attached to  $\mathcal{A}$  (and still far enough from the TuringMobile), which the  
 525 robots avoid when they join  $\mathcal{A}$ . This is because this second area will be used to form the  
 526 pattern.

527 Since  $n$  is known, the TuringMobile knows when it has to interrupt the exploration of the  
 528 plane because all robots have already been found. At this point, the robots switch algorithm:  
 529 one by one, they move to  $\mathcal{B}$  and form the pattern. This task is made possible by the presence  
 530 of the TuringMobile, which gives an implicit order to all robots, and also unambiguously  
 531 defines an embedding of the pattern in  $\mathcal{B}$ . So, each robot is implicitly assigned one point in  
 532  $\mathcal{B}$ , and it moves there when its turn comes.

533 If  $n = 3$  or  $n = 4$ , there are uninteresting ad-hoc algorithms to do this: so, let us assume  
 534 that  $n \geq 5$ . The first to move are the robots in  $\mathcal{A}$ : this part is easy, because they all lie on a  
 535 small segment, which already gives them a total order. The robots only have to be careful  
 536 enough not to collide with other robots before reaching their final positions.

537 When this part is done, there are at least two robots in  $\mathcal{B}$ , all of which have distance  
 538 much smaller than  $d$  from each other. Then the members of the TuringMobile join  $\mathcal{B}$  as well,  
 539 in order from the closest to the farthest. Each of them chooses a position in  $\mathcal{B}$  based on the  
 540 robots already there and the remnants of the TuringMobile. Moreover, the members of the  
 541 TuringMobile that have not started moving to  $\mathcal{B}$  yet cannot be mistaken for robots in  $\mathcal{B}$ ,  
 542 because they are at a greater distance from all others (and vice versa).

543 Note that, when the last robot leaves the TuringMobile and joins  $\mathcal{B}$ , it is able to find its  
 544 final location because there are already at least four robots there, which provide a reference  
 545 frame for the pattern to be formed. When this last robot has taken position in  $\mathcal{B}$ , the pattern  
 546 is formed.

#### 547 4.5 Higher Dimensions

548 Everything we said in this section pertained to robots in the plane. However, we can  
 549 generalize all our results to robots in  $\mathbb{R}^m$ , for  $m \geq 2$ . Recall that, at the end of Section 3.2,  
 550 we have described a TuringMobile for robots in  $\mathbb{R}^m$ , which can move within a specific plane  
 551  $\gamma$  exactly as a bidimensional TuringMobile, but can also move back and forth by  $\mu$  in all  
 552 other directions orthogonal to  $\gamma$ .

553 Now, extending our results to  $\mathbb{R}^m$  actually boils down to exploring the space with a  
 554 TuringMobile: once we can do this, we can easily adapt our techniques for the Near-Gathering

555 and the Pattern Formation problem, with negligible changes.

556 There are several ways a TuringMobile can explore  $\mathbb{R}^m$ : we will only give an example.  
557 Consider the exploration of the plane described at the beginning of this section, and let  $P_i$   
558 be the point reached by the Commander after its  $i$ th move along the spiral-like path depicted  
559 in Figure 3 ( $P_0$  is the initial position of the Commander).

560 Our  $m$ -dimensional TuringMobile starts exploring  $\gamma$  as if it were  $\mathbb{R}^2$ . Whenever it visits  
561 a  $P_i$  for the first time, it goes back to  $P_0$ . From  $P_0$ , it keeps making moves orthogonal to  
562  $\gamma$  until it has seen all points in  $\mathbb{R}^m$  whose projection on  $\gamma$  is  $P_0$  and whose distance from  
563  $P_0$  is at most  $i$ . Then it goes back to  $P_0$ , moves to  $P_1$ , and repeats the same pattern of  
564 moves in the section of  $\mathbb{R}^m$  whose projection on  $\gamma$  is  $P_1$ . It then does the same thing with  
565  $P_2$ , etc. When it reaches  $P_{i+1}$  (for the first time), it goes back to  $P_0$ , and proceeds in the  
566 same fashion. By doing so, it explores the entire space  $\mathbb{R}^m$ .

567 Note that this algorithm only requires the TuringMobile to count how many moves it has  
568 made since the beginning of the execution: thus, the machine only has to memorize a single  
569 integer. The direction of the next move according to the above pattern is then obviously  
570 Turing-computable given the move counter.

## 571 **5** Conclusions

572 We have introduced the TuringMobile as a special configuration of oblivious non-rigid robots  
573 that can simulate a rigid robot with memory. We have also applied the TuringMobile to  
574 some typical robot problems in the context of limited visibility, showing that the assumption  
575 of connectedness of the initial visibility graph can be dropped if a unique TuringMobile is  
576 present in the system. Our results hold not only in the plane, but also in Euclidean spaces  
577 of higher dimensions.

578 The simplest version of the TuringMobile (Section 3.1) consists of only three robots,  
579 and is the smallest possible configuration with these characteristics (Theorems 2 and 3).  
580 Our generalized TuringMobile (Section 3.2), which works in  $\mathbb{R}^m$  and simulates  $k$  registers of  
581 memory, consists of  $3m + 3k$  robots (Theorem 1). We believe we can decrease this number  
582 to  $m + k + 3$  by putting all the Number robots in the same basic component and adopting a  
583 more complicated technique to move them. However, minimizing the number of robots in a  
584 general TuringMobile is left as an open problem.

585 Our basic TuringMobile design works if the robots have the same radius of visibility,  
586 because that allows them to implicitly agree on a unit of distance. We could remove this  
587 assumption and let each of them have a different visibility radius, but we would have to add  
588 a fourth robot to the TuringMobile for it to work (as well as keep the TuringMobile small  
589 compared to *all* these radii).

590 Recall that, in order to encode and decode arbitrary real numbers we used the  $\alpha$  function  
591 and its inverse, which in turn are computed using the arctan and the tan functions. However,  
592 using transcendental functions is not essential: we could achieve a similar result by using  
593 only comparisons and arithmetic operations. The only downside would be that such a real  
594 RAM program would not run in a constant number of machine steps, but in a number of  
595 steps proportional to the value of the number to encode or decode. With this technique, we  
596 would be able to dispense with the trigonometric functions altogether, and have our robots  
597 use only arithmetic operations and square roots to compute their destination points.

## 598 — References

- 599 **1** C. Agathangelou, C. Georgiou, and M. Mavronicolas. A distributed algorithm for gathering  
600 many fat mobile robots in the plane. In *32nd ACM Symposium on Principles of Distributed*  
601 *Computing (PODC)*, 250–259, 2013.
- 602 **2** N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots.  
603 *SIAM Journal on Computing*, 36(1):56–82, 2006.
- 604 **3** A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algo-*  
605 *rithms*. Addison-Wesley, Reading, Massachusetts, 1974.
- 606 **4** H. Ando, Y. Oasa, I. Suzuki, and M. Yamashita. Distributed memoryless point convergence  
607 algorithm for mobile robots with limited visibility. *IEEE Transactions on Robotics and*  
608 *Automation*, 15(5):818–828, 1999.
- 609 **5** B. Bérard, P. Lafourcade, L. Millet, M. Potop-Butucaru, Y. Thierry-Mieg, and S. Tixeuil.  
610 Formal verification of mobile robot protocols. *Distributed Computing*, 29(6):459–487, 2016.
- 611 **6** Q. Bramas and S. Tixeuil. The random bit complexity of mobile robots scattering. *Inter-*  
612 *national Journal of Foundations of Computer Science*, 28(2): 111–134, 2017.
- 613 **7** D. Canepa, X. Défago, T. Izumi, and M. Potop-Butucaru. Flocking with oblivious robots. In  
614 *18th International Symposium on Stabilization, Safety, and Security of Distributed Systems*  
615 *(SSS)*, 94–108, 2016.
- 616 **8** S. Cicerone, G. Di Stefano, and A. Navarra. Asynchronous pattern formation: the effects  
617 of a rigorous approach. *arXiv:1706.02474*, 2017.
- 618 **9** S. Cicerone, G. Di Stefano, and A. Navarra. Minimum-traveled-distance gathering of obliv-  
619 ious robots over given meeting points. In *10th International Symposium on Algorithms and*  
620 *Experiments for Sensor Systems (Algosensors)*, 57–72, 2014.
- 621 **10** M. Cieliebak, P. Flocchini, G. Prencipe, and N. Santoro. Distributed computing by mobile  
622 robots: Gathering. *SIAM Journal on Computing*, 41(4):829–879, 2012.
- 623 **11** R. Cohen and D. Peleg. Convergence properties of the gravitational algorithm in asyn-  
624 chronous robot systems. *SIAM Journal on Computing*, 34(6):1516–1528, 2005.
- 625 **12** P. Courtieu, L. Rieg, S. Tixeuil, and X. Urbain. Certified universal gathering in  $\mathbb{R}^2$  for obliv-  
626 ious mobile robots. In *30th International Symposium on Distributed Computing (DISC)*,  
627 187–200, 2016.
- 628 **13** P. Courtieu, L. Rieg, S. Tixeuil, and X. Urbain. Impossibility of gathering, a certification.  
629 In *Information Processing Letters*, 115(3):447–452, 2015.
- 630 **14** S. Das, P. Flocchini, N. Santoro, and M. Yamashita. Forming sequences of geometric  
631 patterns with oblivious mobile robots. *Distributed Computing*, 28(2):131–145, 2015.
- 632 **15** X. Défago, M. Gradinariu, S. Messika, P. Raipin-Parvédy, and S. Dolev. Fault-tolerant and  
633 self-stabilizing mobile robots gathering. In *20th International Symposium on Distributed*  
634 *Computing (DISC)*, 46–60, 2006.
- 635 **16** B. Degener, B. Kempkes, P. Kling, and F. Meyer auf der Heide. Linear and competi-  
636 tive strategies for continuous robot formation problems. *ACM Transactions on Parallel*  
637 *Computing*, 2(1): 2:1–2:18, 2015.
- 638 **17** B. Degener, B. Kempkes, T. Langner, F. Meyer auf der Heide, P. Pietrzyk, and R. Wat-  
639 tenhofer. A tight runtime bound for synchronous gathering of autonomous robots with  
640 limited visibility. In *23rd ACM Symposium on Parallelism in Algorithms and Architectures*  
641 *(SPAA)*, 139–148, 2011.
- 642 **18** G. Di Luna, P. Flocchini, N. Santoro, and G. Viglietta. TuringMobile: a turing machine  
643 of oblivious mobile robots with limited visibility and its applications. *arXiv:1709.08800*,  
644 2017.
- 645 **19** P. Flocchini, G. Prencipe, and N. Santoro. *Distributed Computing by Oblivious Mobile*  
646 *Robots*. Morgan & Claypool, 2012.



- 647 **20** P. Flocchini, G. Prencipe, N. Santoro, and G. Viglietta. Distributed computing by  
648 mobile robots: Uniform circle formation. *Distributed Computing*, 2016, to appear,  
649 doi:10.1007/s00446-016-0291-x.
- 650 **21** P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer. Gathering of asynchronous robots  
651 with limited visibility. *Theoretical Computer Science*, 337(1–3):147–168, 2005.
- 652 **22** P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer. Arbitrary pattern formation by  
653 asynchronous, anonymous, oblivious robots. *Theoretical Computer Science*, 407(1–3):412–  
654 447, 2008.
- 655 **23** N. Fujinaga, Y. Yamauchi, S. Kijima, and M. Yamashita. Pattern formation by oblivious  
656 asynchronous mobile robots. *SIAM Journal on Computing*, 44(3):740–785, 2015.
- 657 **24** T. Izumi, M. Gradinariu Potop-Butucaru, and S. Tixeuil. Connectivity-preserving scatter-  
658 ing of mobile robots with limited visibility In *12th International Symposium on Stabiliza-  
659 tion, Safety, and Security of Distributed Systems (SSS)*, 319–331, 2010.
- 660 **25** L. Pagli, G. Prencipe, and G. Viglietta. Getting close without touching: Near-Gathering  
661 for autonomous mobile robots. *Distributed Computing*, 28(5):333–349, 2015.
- 662 **26** F.P. Preparata and M.I. Shamos. *Computational Geometry*. Springer-Verlag, Berlin and  
663 New York, 1985.
- 664 **27** H. Rogers, Jr. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill,  
665 1967.
- 666 **28** M.I. Shamos. *Computational Geometry*. Ph.D. thesis, Department of Computer Science,  
667 Yale University, 1978.
- 668 **29** I. Suzuki and M. Yamashita. Distributed anonymous mobile robots: formation of geometric  
669 patterns. *SIAM Journal on Computing*, 28(4):1347–1363, 1999.
- 670 **30** M. Yamashita and I. Suzuki. Characterizing geometric patterns formable by oblivious  
671 anonymous mobile robots. *Theoretical Computer Science*, 411(26–28):2433–2453, 2010.
- 672 **31** Y. Yamauchi, T. Uehara, S. Kijima, and M. Yamashita. Plane formation by synchronous  
673 mobile robots in the three-dimensional Euclidean space. *Journal of the ACM* 64(3): 16:1–  
674 16:43, 2017.