Fault-Induced Dynamics of Oblivious Robots on a Line ⋆

Jean-Lou De Carufel and Paola Flocchini

University of Ottawa, {jdecaruf,pflocchi}@uottawa.ca

Abstract. The study of computing in presence of faulty robots in the LOOK-COMPUTE-MOVE model has been the object of extensive investigation, typically with the goal of designing algorithms tolerant to as many faults as possible. In this paper, we initiate a new line of investigation on the presence of faults, focusing on a rather different issue. We are interested in understanding the dynamics of a group of robots when they execute an algorithm designed for a fault-free environment, in presence of some undetectable crashed robots. We start this investigation focusing on the classic point-convergence algorithm by Ando et al. [2] for robots with limited visibility, in a simple setting (which already presents serious challenges): the robots operate fully synchronously on a line, and at most two of them are faulty. Interestingly, and perhaps surprisingly, the presence of faults induces the robots to perform some form of *scattering*, rather than *point-convergence*. In fact, we discover that they arrange themselves inside the segment delimited by the two faults in interleaved sequences of equidistant robots.

1 Introduction

Consider a group of robots represented as points, which operate in a continuous space according to the LOOK-COMPUTE-MOVE model [16]: when active, a robot LOOKS the environment obtaining a snapshot of the positions of the other visible robots, it COMPUTES a destination point on the basis of such a snapshot, and it MOVES there. As typically assumed by the model, the robots are *anonymous* (i.e., they are identical), *autonomous* (without central or external control), *oblivious* (they have no memory of past activations), *disoriented* (they do not agree on a common coordinate systems), *silent* (they have no means of explicit communication). These systems of autonomous robots have been extensively investigated under different assumptions on the various model parameters (different levels of synchrony, level of agreement on the coordinate system, etc.), and most algorithms in the literature are designed for fault-free groups of robots (e.g., see [7, 12–15, 17, 19–21]).

⋆ This work has been supported in part by the Natural Sciences and Engineering Research Council of Canada through the Discovery Grant program; by Prof. Flocchini's University Research Chair.

There are several studies that consider the presence of faults: crashes (robots that are never activated) or byzantine (robots that behave differently than intended). The goal, in these cases, has been to design fault-tolerant algorithms focusing on the maximum amount of faults that can be tolerated for a solution to exist in a given model (e.g., see [1, 3–6, 11]). For a detailed account of the current investigations see [11].

In this paper, we consider a rather different question in presence of faulty robots that has never been asked before. Given an algorithm designed to achieve a certain global goal by a group of fault-free robots, what is the behaviour of the robots in presence of crash faults ? Clearly, in most cases, the original goal is not achieved, but the theoretical interest is in characterizing the dynamics of the non-faulty robots induced by the presence of the faulty ones, from arbitrary initial configurations. Apart from the theoretical curiosity, this approach can be seen as a first step toward the study of the interaction between heterogeneous groups of robots operating in the same space, each following a different algorithm. In fact, the dynamics resulting from the presence of different teams following different and possibly conflicting rules in the environment is an important area of investigation that has never been studied.

We start this new line of investigation focusing on the classic point-convergence algorithm by Ando et al. [2] for robots with limited visibility, and considering one of the simplest possible settings, which already proves to be challenging: fully synchronous robots (FSynch) moving in a 1-dimensional space (a line), in presence of at most two faults. In a line, the convergence algorithm prescribes each robot to move to the center of the leftmost and rightmost visible robots and, in absence of faults, starting from a configuration where the robots' "visibility graph" is connected, the robots are guaranteed to converge toward a point. It is not difficult to see that with a single fault, the robots successfully converge toward the faulty robot. The presence of multiple faults, however, gives rise to intricate dynamics, and the analysis of the robots behavior is already quite complex with just two. The case of more than two faults is left for further study.

Interestingly, and perhaps surprisingly, the presence of faults induces the robots to perform some form of *scattering*, rather than *gathering*. In fact, we prove that they arrange themselves inside the segment delimited by the two faults in interleaved sequences composed of equidistant robots. The structure that they form has a hierarchical nature: robots organize themselves in groups where a group of some level converges to an equidistant distribution between the first and the last robots of that group. Moreover, the first and the last robots of that group belong to a lower level group. Also interesting to note is the rather different dynamics that arises when moving to the middle between two robots, depending on the choice of the robots: when considering the *closest* neighbours, the result is an equidistant distribution (scattering algorithm of [13]), when instead selecting the *leftmost* and *rightmost* visible robots the result is a more complex structure of sequences of robots, each converging to an equidistant distribution. The main difficulty of our analysis is to show that the robots indeed form this special combination of sequences: the convergence of each sequence is then derived from a generalization of the result by [13].

Finally observe that the 2-dimensional case has a rather different nature. In fact, in contrast to the 1-dimensional setting, where any initial configuration converges toward a pattern, when robots move on the plane oscillations are possible, even with just two faults. The investigation of this case is left for future study.

Due to lack of space, most proofs are only sketched. The full version of the paper can be found in [10].

2 Preliminaries

2.1 Model and Notation

Let X denote a set of identical point-form robots moving on a line, simultaneously activated in synchronous time steps according to the LOOK-COMPUTE-MOVE model [16]. The robots have limited visibility. In the LOOK phase, they "see" the positions of the robots within their visibility radius V, then they all COMPUTE a destination point, and they MOVE to that point. The robots are oblivious in the sense that the computation at time t solely depends on the positions of the robots perceived at that step. We assume that two robots, arbitrarily placed, are permanently faulty (i.e., they are stationary and inactive). Their faulty status, however, is not visible and they appear identical to the others. Let $X(t) = \{x_0(t), x_1(t), ..., x_n(t)\}$ be the set of robots at time t. Let x denote a robot $x \in X$ and $x(t)$ its position at time t with respect to the leftmost faulty robot. With an abuse of notation $x(t)$ may indicate both the robot itself and its position at time t. Robots do not necessarily occupy distinct positions. For instance we might have $x_i(t) = x_j(t)$ where $0 \leq i, j \leq n$ are two different indices. Note, however, that non-faulty robots in the same position behave in the same way and can be considered as a single one. Indeed, when non-faulty robots end up in the same position, we say that they "merge" and from that moment on they will be considered as one.

We denote the distance between robots x and y at time t by $|x(t) - y(t)|$. We denote by $[\alpha, \beta]$ the interval of real numbers starting at $\alpha \in \mathbb{R}$ and ending at $\beta \in \mathbb{R}$, where $\alpha \leq \beta$. Let $N(x(t))$ be the set of robots visible by x at time t, that is: $N(x(t))$ is the set of robots y such that $|x(t) - y(t)| \leq V$. Let $r(x(t))$ (resp. $l(x(t))$) denote the rightmost (resp. the leftmost) robot visible by x at time t. If no robot is visible to the right (resp. to the left), then $r(x(t)) = x(t)$ (resp. $l(x(t)) = x(t)$). We say that a configuration of robots $X = \{x_0, x_1, ..., x_n\}$ converges to a pattern $P = \{p_0, p_1, ..., p_n\}$ if for all $0 \leq i \leq n$, $x_i(t) \rightarrow p_i$ as $t \rightarrow \infty$.

2.2 Background Results: Point-Convergence and Scattering

Point-convergence [2]. A classical problem for oblivious robots is *point-convergence*: the robots, initially placed in arbitrary positions, must converge toward the same point, not established a-priori. A solution to this problem is given by the well known algorithm by Ando et al. [2]. The algorithm achieves convergence to a point, not only in synchronous systems, but also when at each time step, only a subset of the robots is activated (semi-synchronous scheduler SSynch), as long as every robot is activated infinitely often. The robots are initially placed in arbitrary positions in a 2-dimensional space and have limited visibility. The algorithm prescribes each robot to move toward the centre of the smallest enclosing circle that contains all the robots up to a certain distance, guaranteeing any pair of robots to maintain visibility in spite of each others possible movement.

When the space where the robots can move is a line, the algorithm (Convergence1D) becomes quite simple because the smallest enclosing circle of the visible robots is the segment delimited by the leftmost and rightmost visible robots, and a robot moves to occupy the mid-point between them.

Theorem 1. *[2] Executing Algorithm* Convergence1D *in* FSynch *or* SSynch*, the robots converge to a point.*

Scattering on a segment [9]. In [9], a classical *scattering* algorithm for robots in 1-dimensional systems has been analyzed both in FSynch and SSynch. A variant of this result (Theorem 3) will be heavily used in this paper. We briefly describe the main result and its generalization.

Consider a set of oblivious robots $X = \{x_0, x_1, ..., x_n\}$ on a line, where x_0 and x_n do not move (equivalently, this can be considered as a segment delimited by the positions of x_0 and x_n). Let $D = |x_n(0) - x_0(0)|$. In [9], the robots are assumed to be able to see the closest robot on each side, while x_0 and x_n know they are the delimiters of the segment. The algorithm of [9] (Spreading) makes the robot converge to a configuration where the distance between consecutive robots tends to $\frac{D}{n}$ by having the extremal robots never move and the others move to the middle point between the two neighbours.

Theorem 2. *[9] Executing Algorithm* Spreading *in* FSynch *or in* SSynch *on the set of robots R where the first and the last robots do not move, the robots converge to equidistant positions.*

The theorem can be generalized in FSynch to the case when x_0 and x_n are not stationary, but are each converging toward a point (resp. x'_0 and x'_n). The proof is technical, but it essentially follows the same lines of the proof of [9], and can be found in the full version of the paper [10].

Theorem 3. *Let $X = \{x_0, x_1, ..., x_n\}$ where $x_0(t) \to x'_0$ and $x_n(t) \to x'_n$ as $t \to \infty$. Executing Algorithm* Spreading *in* FSynch *on robots $\{x_1, ..., x_{n-1}\}$, the robots converge to equidistant positions between x'_0 and x'_n.*

3 Robots' Dynamics in Presence of Two Faults

It is not difficult to see that, if the configuration contains one faulty robot, the other robots converge toward it. We then focus on the case when the system contains two faults and we show that, starting from an arbitrary configuration, the system converges towards a limit configuration.

For the rest of this paper, we will always denote by x_0 (resp. by x_n) the leftmost (resp. the rightmost) faulty robot. Moreover, for simplicity, x_0 is considered to be at position 0 (note that there could be robots initially placed in negative positions).

3.1 Basic Properties

We start with a series of lemmas leading to the proof of two crucial properties: there exists a time after which the robots preserve their farthest neighbours (Theorem 4) and when the number of different positions occupied by them becomes constant (Corollary 1).

Lemma 1 (No Crossing). *If x and z are two non-faulty robots and $x(t) < z(t)$, then $x(t+1) \leq z(t+1)$.*

Proof. Since $x(t) < z(t)$, we have that $r(x(t)) \leq r(z(t))$ and $l(x(t)) \leq l(z(t))$ by definition. It follows that $x(t+1) = \frac{l(x(t))+r(x(t))}{2} \leq \frac{l(z(t))+r(z(t))}{2} = z(t+1)$. □

With the next two lemmas we show that all robots, except possibly two, eventually enter the segment $[x_0, x_n]$ delimited by the two faulty robots. At most two robots might perpetually stay outside of it, one to the left of x_0 and one to the right of x_n. If this is the case, however, the two outsiders converge to x_0 and x_n, respectively.

Lemma 2 (No More Crossing). *If x is a non-faulty robot, it will cross at most a finite number of times with a faulty robot.*

Proof. (Sketch) Using Lemma 1, we can show that there is a non-faulty robot x_ℓ (resp. x_r) that will stay the leftmost (resp. the rightmost) non-faulty robot for all $t \geq 0$.

We first consider the faulty robot x_0. If $x_\ell(t) \in [x_0, x_n]$ for some time t, then $l(x_\ell(t)) = x_0$, from which $x_\ell(t') \in [x_0, x_n]$ for all $t' \geq t$. Otherwise, for all $t \geq 0$, we have $x_\ell(t) < x_0$, $l(x_\ell(t)) = x_\ell(t)$ and $r(x_\ell(t)) > x_\ell(t)$. Thus, $x_\ell(t)$ is strictly increasing as $t \to \infty$. Therefore, $x_\ell(t) \to x^*$ as $t \to \infty$, for some $x^* \leq x_0$. We can prove that $x^* = x_0$ by showing that all robots which are to the left of x_0 are attracted by x_0. Then, we can prove that all non-faulty robots in the interval $[x_\ell(t), x_0]$ will merge with x_ℓ after a finite number of steps. Therefore, all non-faulty robots in the interval $[x_\ell(t), x_0]$ will cross at most a finite number of times with x_0.

A symmetric argument for x_n completes the proof. □

Lemma 3. *There is a time $t \geq 0$ such that either one of the following two scenarios happens:*
- All robots are inside the line segment $[x_0, x_n]$ and will stay there for all $t' \geq t$.
- All robots, except for at most two of them (x_ℓ and x_r), are inside $[x_0, x_n]$ and will stay there for all $t' \geq t$. We have that $x_\ell(t') < x_0$ and $x_r(t') > x_n$ for all $t' \geq t$. Moreover, $x_\ell(t) \to x_0$ and $x_r(t) \to x_n$ as $t \to \infty$.

Proof. (Sketch) The proof is similar to the one of Lemma 2. After a finite number of steps, all non-faulty robots in the interval $[x_\ell(t), x_0]$ will merge with x_ℓ. Hence, after a finite number of steps, there is only one robot remaining to the left of x_0 (two robots merging together are considered as a single robot). A symmetric argument holds for x_r. □

The two dissident robots from the previous lemma are called *outsiders*. Since $x_\ell(t') < x_0$ and $x_r(t') > x_n$ for all $t' \geq t$, and since $x_\ell(t') \to x_0$ and $x_r(t') \to x_n$ as $t \to \infty$, we can ignore them without loss of generality. For the rest of the paper, we suppose that all robots are inside $[x_0, x_n]$ and will stay there for all $t' \geq t$.

We now show that during the evolution of the system, a robot never loses visibility of the robots seen in the past.

Lemma 4 (Preserved Visibility). *Let $y \in N(x(t))$. For all $t' > t$, $y \in N(x(t'))$.*

Proof. (Sketch) Let $y \in N(x(t))$. Without loss of generality, $y(t)$ is to the left of $x(t)$, from which $0 < x(t) - y(t) \leq V$. If both x and y are faulty, they do not move and the result follows. Otherwise, we write $x(t+1) - y(t+1) = \frac{l(x(t))+r(x(t))}{2} - \frac{l(y(t))+r(y(t))}{2}$, which can be shown to be upper bounded by V. □

During the execution of the algorithm, robots could cross each other (*crossing*), they could merge and occupy the same position (*merging*), and could enter the visibility range of a robot (*inclusion*). A size-stable time is when inclusions, crossings and mergings cease to happen and all robots are inside the segment.

Definition 1 (Size-Stable Time). *A time t_0 is called a* size-stable time *if: for all $t \geq t_0$, there are no inclusions, mergings or crossings in the system, and at most one agent stays permanently on each side of the line segment $[x_0, x_n]$ converging toward x_0 and x_n, respectively.*

From Lemmas 1 and 2, after a finite number of steps, no two robots are *crossing* each others. From Lemma 3, either all robots are inside the line segment $[x_0, x_n]$ after a finite number of steps, or at most two robots will stay outside of the line segment $[x_0, x_n]$ for all time $t \geq 0$. We then get the following corollary.

Corollary 1. *For all set of robots X, there exists a size-stable time t_0.*

Finally, from Lemmas 1, 2 and 4, and Corollary 1, we can conclude that at any time after a size-stable time t is reached, the farthest left and right neighbours, namely $l(x(t))$ and $r(x(t))$, of any robot x will never change.

Theorem 4 (Preserved-farthest-neighbours). *Let t be a size-stable time and $x \in \mathcal{R}$ be a robot. For all $t' \geq t$, $r(x(t')) = r(x(t))$ and $l(x(t')) = l(x(t))$.*

For the rest of the paper, we suppose that the earliest size-stable time is 0. Thus, from Corollary 1, for all $t \geq 0$, t is a size-stable time.

3.2 Convergence of Mutual Chains

We now define the notion of *mutual chain* as a set of robots that are mutually the farthest from each other.

Definition 2 (Mutual Chain). *Let $0 \leq k \leq n$ be an integer and $t \geq 0$ be any size-stable time. A* mutual chain at time t *(or* mutual chain *for short) is a configuration $C(t) = \{x'_1(t), x'_2(t), ..., x'_k(t)\} \subset X(t)$ made of k robots such that for all $1 \leq i \leq k-1$, $l(x'_{i+1}(t)) = x'_i(t)$ and $r(x'_i(t)) = x'_{i+1}(t)$ (refer to Figure 1).*

If $r(x_i(t)) = x_j(t)$ and $l(x_j(t)) = x_i(t)$, we say that x_i and x_j are mutually chained at time t *or that $x_i(t)$ and $x_j(t)$ are* mutually chained.

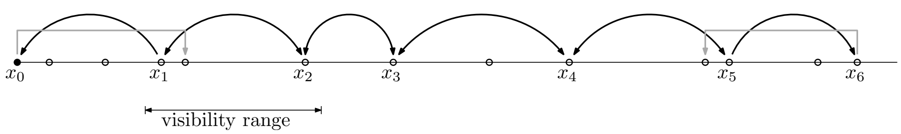

Fig. 1. A mutual chain of robots $C(t) = \{x_1(t), x_2(t), x_3(t), x_4(t), x_5(t)\}$ anchored in x_0 and x_6, where the arrows indicate farthest visibility.

The *anchors* of a mutual chain $C(t) = \{x'_1(t), x_2(t), ..., x'_k(t)\}$ are the farthest left neighbour of $x'_1(t)$ and the farthest right neighbour of $x'_k(t)$.

Definition 3 (Anchors). *Given a mutual chain $C(t) = \{x'_1(t), x'_2(t), \ldots, x'_k(t)\}$, we say that $l(x'_1(t))$ and $r(x'_k(t))$ are the left and right* anchors *of $C(t)$ (or that $C(t)$ is* anchored *at $l(x'_1(t))$ and $r(x'_k(t))$) (refer to Figure 1).*

Note that the definition of anchor allows the anchors of a mutual chain to be part of the mutual chain (refer to Figure 2). The anchors do not have to be

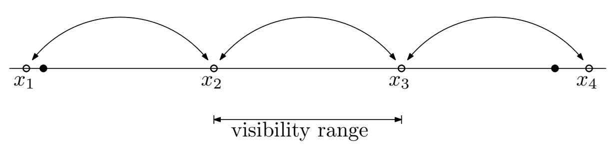

Fig. 2. Configuration $\{x_1, x_2, x_3, x_4\}$ is a mutual chain. It is anchored at x_1 and x_4.

faulty robots for this situation to happen. Moreover, the definition of mutual chain allows a mutual chain to possibly contain only one robot. Indeed, any robot x forms a mutual chain $\{x(t)\}$ anchored at $l(x(t))$ and $r(x(t))$.

We now prove the formation, during the execution of the algorithm, of a special unique mutual chain called *primary chain*. Intuitively, the primary chain is a mutual chain starting from x_0 and ending in x_n. We will then introduce a hierarchical notion of mutual chains with different levels, where chains of some level are anchored in lower level ones. Moreover, we will show that the robots will eventually arrange themselves in such a hierarchical structure of mutual chains.

Theorem 5 (Primary Chain). *There exists a configuration of robots* $\mathcal{C}_1 = \{x'_0, x'_1, x'_2, ..., x'_k\} \subseteq X$ *such that at any size-stable time* $t > 0$, $\mathcal{C}_1(t)$ *is a mutual chain anchored at* x_0 *and* x_n*, where* $x'_0 = x_0$ *and* $x'_k = x_n$*. This mutual chain is called the* primary chain *of* X *and it is unique.*

Before we prove Theorem 5, we need the following technical lemma (whose proof can be found in the full version of the paper [10]). Intuitively, when the distance between two mutually chained robots tends to V (as $t \to \infty$), this limit behaviour propagates to the leftmost and rightmost visible robots.

Lemma 5. *Let* $x'_{\alpha+1}, x'_{\alpha+2} \in X$ *such that for all* $t \geq 0$*:*
- $x'_{\alpha+1}(t)$ *and* $x'_{\alpha+2}(t)$ *are mutually chained,*
- $d(t) = x'_{\alpha+2}(t) - x'_{\alpha+1}(t) \to V$*, as* $t \to \infty$
- $l(x'_{\alpha+1}(t)) \neq x'_{\alpha+1}(t)$
- $r(x'_{\alpha+2}(t)) \neq x'_{\alpha+2}(t)$.

Then, $r(x'_{\alpha+2}(t)) - x'_{\alpha+2}(t) \to V$ *and* $x'_{\alpha+1}(t) - l(x'_{\alpha+1}(t)) \to V$*, as* $t \to \infty$.

Proof. (Sketch) The robots $x'_{\alpha+1}(t)$ and $x'_{\alpha+2}(t)$ are mutually chained and $x'_{\alpha+2}(t) - x'_{\alpha+1}(t) \to V$, as $t \to \infty$. Since $x'_{\alpha+1}(t+1)$ always places itself in the middle of $l(x'_{\alpha+1}(t))$ and $r(x'_{\alpha+1}(t)) = x'_{\alpha+2}(t)$, we must have that $x'_{\alpha+1}(t) - l(x'_{\alpha+1}(t)) \to V$, as $t \to \infty$. The same reasoning applies for $r(x'_{\alpha+2}(t))$ and $x'_{\alpha+2}(t)$. This can be formalized using the formal definition of limits. □

Proof. (Theorem 5)

[Uniqueness] We first explain that if the primary chain exists, then it is unique. Since $x_0 = x'_0$ and $x_n = x'_k$ are part of the mutual chain, starting at x_0, we get $x'_1 = r(x_0)$ and $x'_{i+1} = r(x'_i)$ for all $0 \leq i \leq k-1$, where $x'_k = x_n$. So each x'_i is uniquely defined.

[Existence] We prove the existence of the primary chain by contradiction. Let us summarize the steps of the proof. We assume that there does not exist any mutual chain. 1) We construct a particular configuration, composed by a forward-chain from x_0 connecting each node to its farthest right neighbour until x_n is reached and a backward chain from x_n connecting each node to its farthest left neighbour back to x_0. 2) We then show that the two chains converge to each other, i.e., they converge to a single chain, called *right-left chain.* This construction does not directly guarantee that the right-left chain is a mutual chain. We then show a contradiction, reasoning on the total length of the segment delimited by x_0 and x_n. 3) A consequence of the right-left chain not being a mutual chain is that the total length of the segment between x_0 and x_n is strictly smaller than $(j+1)V$ (where $j+1$ is the number of intervals between consecutive robots in the chain). 4) On the other hand, each such interval converges to V, thus implying that the total length of the segment is a number arbitrarily close to $(j+1)V$ (by Lemma 5). The contradiction implies that the right-left chain is indeed mutual.

1) Construction of forward and backward chains. Let us consider a configuration of robots $\{x'_0(t), x'_1(t), ..., x'_{j+1}(t)\} \subseteq X(t)$, called *forward chain* (refer to Figure 3), such that: $x'_0(t) = x'_0 = x_0$, $x'_{i+1}(t) = r(x'_i(t))$ for all $0 \le i \le j < n$, and $x'_{j+1}(t) = x'_{j+1} = x_n$.

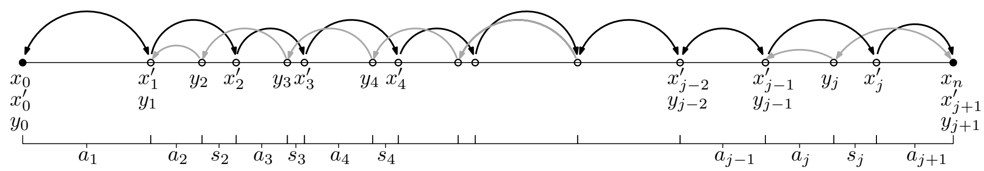

Fig. 3. Illustration of the proof of Theorem 5.

We define another configuration of robots, called *backward chain*, $\{y_0(t), y_1(t), \ldots, y_{j+1}(t)\} \subseteq X(t)$ as follows. Let $y_{j+1}(t) = x'_{j+1}(t)$ and for all $0 \le i \le j$, let $y_i(t) = l(y_{i+1}(t))$ (refer to Figure 3). Let us call the union of the two chains *right-left chain*. We can prove two useful properties about the right-left chain.

Property 1 (Alternation.) *For all* $1 \le i \le j+1$, $x'_{i-1}(t) < y_i(t) \le x'_i(t)$.
Property 2 (Starting point.) *We have that* $y_0(t) = y_0 = x_0$.

2) Convergence of forward and backward chains to a right-left chain. Notice that since the forward chain $\{x'_0(t), x'_1(t), ..., x'_{j+1}(t)\}$ is not a mutual chain, there exists an i with $1 \le i \le j$ such that $x'_{i-1}(t) < y_i(t) < x'_i(t)$. For all $1 \le i \le j+1$, let $a_i(t) = y_i(t) - x'_{i-1}(t)$ and $s_i(t) = x'_i(t) - y_i(t)$. Our aim, in the following, is to prove that $x'_i(t)$ and $y_i(t)$ get arbitrarily close when $t \to \infty$. From Property 1, we have $a_i(t) > 0$ and $s_i(t) \ge 0$ for all $1 \le i \le j+1$. Moreover, $s_i(t) = 0$ if and only if $y_i(t) = x'_i(t)$. Notice that $l(x'_i(t-1)) \le x'_{i-1}(t-1)$, otherwise there would be a contradiction with the fact that $r(x'_{i-1}(t-1)) = x'_i(t-1)$. Therefore, $x'_i(t) = \frac{l(x'_i(t-1)) + r(x'_i(t-1))}{2} \le \frac{x'_{i-1}(t-1) + x'_{i+1}(t-1)}{2}$, from which $x'_0(t) \le 0$, $x'_{j+1}(t) \le x_n$ and

$$x'_i(t) \le x'_{i-1}(t-1) + \frac{1}{2}(a_i(t-1) + s_i(t-1) + a_{i+1}(t-1) + s_{i+1}(t-1)) \quad (1)$$

for all $1 \le i \le j$. Moreover, notice that $r(y_i(t-1)) \ge y_{i+1}(t-1)$, otherwise there would be a contradiction with the fact that $l(y_{i+1}(t-1)) = y_i(t-1)$. Therefore, $y_i(t) = \frac{l(y_i(t-1)) + r(y_i(t-1))}{2} \ge \frac{y_{i-1}(t-1) + y_{i+1}(t-1)}{2}$, from which $y_0(t) \ge 0$, $y_{j+1}(t) \ge x_n$ and

$$y_i(t) \ge y_{i-1}(t-1) + \frac{1}{2}(s_{i-1}(t-1) + a_i(t-1) + s_i(t-1) + a_{i+1}(t-1)) \quad (2)$$

for all $1 \le i \le j$. Since $s_i(t) = x'_i(t) - y_i(t)$, by subtracting (2) from (1) we obtain $s_0(t) \le 0$, $s_{j+1}(t) \le 0$ and

$$s_i(t) \le \frac{1}{2}(s_{i-1}(t-1) + s_{i+1}(t-1)) \quad (3)$$

for all $1 \leq i \leq j$. We are now ready to prove that for all $0 \leq i \leq j+1$, $s_i(t) \to 0$ as $t \to \infty$, implying that $y_i(t) \to x'_i(t)$ as $t \to \infty$. Notice that we already have $y_0(t) = x'_0(t)$ and $y_{j+1}(t) = x'_{j+1}(t)$ by definition. By unfolding (3), we get

$$s_i(t) \leq \frac{1}{2^t} \sum_{k=0}^{t} \binom{t}{k} s_{i-t+2k}(0),$$

where $s_i(t) = 0$ for all $i \leq 0$ and $i \geq j+1$.

In order to determine the limit of $s_i(t)$ when $t \to \infty$, we need to make a few observations. First of all, the $s_i(t)$'s in the summation with $i \leq 0$ or $i \geq j+1$ are all equal to zero. In other words, regardless of the value of t, there are at most j non-zero values in the summation. These j values correspond to the j-central binomial coefficients. Also note that since the segment delimited by the two faulty robots has a constant size, the values of the s_i's are bounded. Let C be the value of the largest such s_i ever occurring. Since the largest binomial coefficient is the central one (or the central ones for odd values of t), we can write $0 \leq s_i(t) \leq \frac{1}{2^t} j \binom{t}{\lfloor \frac{t}{2} \rfloor} C$. Since[1] $\binom{t}{\lfloor \frac{t}{2} \rfloor} \sim \frac{2^t}{\sqrt{\pi \frac{t}{2}}}$, we have

$$0 \leq \lim_{t \to \infty} s_i(t) \leq \lim_{t \to \infty} \frac{1}{2^t} j \binom{t}{\lfloor \frac{t}{2} \rfloor} C = \lim_{t \to \infty} \frac{1}{2^t} j \frac{2^t}{\sqrt{\pi \frac{t}{2}}} C = 0,$$

from which $\lim_{t \to \infty} s_i(t) = 0$. We are now ready to derive a contradiction.

3) Length of the segment strictly smaller than $(j+1)V$. Since the right-left chain is not a mutual chain, and x_0 and x_n are not moving, the distance between x_0 and x_n must be strictly smaller than $(j+1)V$ (otherwise x'_j and y_j would necessarily coincide, for all j). So, there exists a real number $\delta > 0$ such that $x_n - x_0 = (j+1)V - \delta$.

4) Distance between $x'_i(t)$ and $x'_{i+1}(t)$ tending to V. Let us consider any sub-chain of the right-left chain for which the x'_i and the y_i are distinct except for the extremal ones. More precisely, let α and β be two indices such that $x'_\alpha = y_\alpha$, $x'_\beta = y_\beta$ and $x'_i \neq y_i$ for all $\alpha < i < \beta$ (refer to Figure 4). Notice that $l(x'_{\alpha+1}) = x'_\alpha$, otherwise this would contradict the fact that $l(y_{\alpha+1}) = x'_\alpha$. We also have $r(y_{\beta-1}) = x'_\beta$, otherwise this would contradict the fact that $r(x'_{\beta-1}) = x'_\beta$. Therefore, $l(x'_{\alpha+1}) = x'_\alpha$, $r(x'_{\alpha+1}) = x'_{\alpha+2}$, $l(y_{\beta-1}) = y_{\beta-2}$ and $r(y_{\beta-1}) = y_\beta = x'_\beta$. This implies that $k \geq i+3$, otherwise $x'_{\alpha+1}$ and $y_{\beta-1}$ would have the same leftmost and rightmost visible robots and they would merge in one step, which is not possible at a size-stable time. Since there cannot be any merging, given that $l(y_{\alpha+1}) = y_\alpha = x'_\alpha$, we must also have that $x'_{\alpha+2}$ is not visible from $y_{\alpha+1}$ at any time. Therefore, for all $t \geq 0$, $s_{\alpha+1}(t) + a_{\alpha+2}(t) + s_{\alpha+2}(t) > V$. Since $r(x'_{\alpha+1}) = x'_{\alpha+2}$, for all $t \geq 0$, $a_{\alpha+2}(t) + s_{\alpha+2}(t) \leq V$. Together with the fact that $s_{\alpha+1}(t) \to 0$ and $s_{\alpha+2}(t) \to 0$ as $t \to \infty$, we get that $a_{\alpha+2}(t) \to V$ as $t \to \infty$.

[1] We write $f(t) \sim g(t)$ whenever $\lim_{t \to \infty} \frac{f(t)}{g(t)} = 1$.

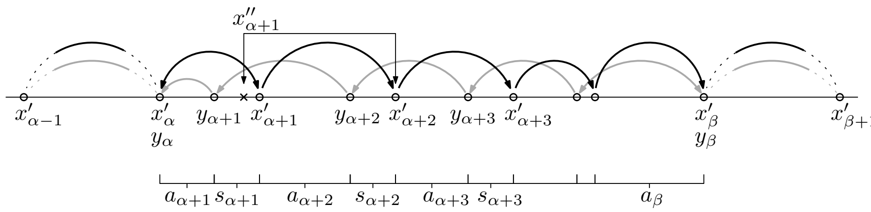

Fig. 4. Illustration of the contradiction in the proof of Theorem 5 (propagation of distance V). We do not make any assumption about $x'_{\alpha-1}$ being equal or not to $y_{\alpha-1}$, nor about $x'_{\beta+1}$ being equal or not to $y_{\beta+1}$.

Therefore, $x'_{\alpha+2}(t) - x'_{\alpha+1}(t) \to V$ as $t \to \infty$. Our goal is to apply Lemma 5 and conclude that $x'_{\alpha+1}(t) - x'_{\alpha} \to V$ and $x'_{\alpha+3}(t) - x'_{\alpha+2} \to V$ as $t \to \infty$. However, since $x'_{\alpha+1}(t)$ and $x'_{\alpha+2}(t)$ are not mutual, we cannot apply the lemma directly. Due to lack of space, we only sketch the idea to circumvent this problem (refer to [10] for full details). We can prove that there is a robot $x''_{\alpha+1}(t)$, satisfying $y_{\alpha+1}(t) \le x''_{\alpha+1}(t) \le x'_{\alpha+1}(t)$, that is mutually chained with $x'_{\alpha+2}(t)$. Intuitively, since $y_{\alpha+1}(t) \to x'_{\alpha+1}(t)$ as $t \to \infty$, and since $x''_{\alpha+1}(t) \in [y_{\alpha+1}(t), x'_{\alpha+1}(t)]$, $x''_{\alpha+1}$ behaves the same way $x'_{\alpha+1}$ does. But since $x''_{\alpha+1}(t)$ is mutually chained with $x'_{\alpha+2}(t)$, we can apply Lemma 5. We can repeat the same argument and show that this propagates to all x'_i's, from which we get that for all $0 \le i \le j$, $x'_{i+1}(t) - x'_i \to V$ as $t \to \infty$. Therefore, the total distance between x_0 and x_n is arbitrarily close to $(j+1)V$. This contradicts the fact that $x_n - x_0 = (j+1)V - \delta$. □

In the proof of Theorem 5, we showed the existence of a unique mutual chain called the primary chain. Intuitively, we say that a configuration of robots is a secondary chain if it is a mutual chain anchored at two robots that belong to the primary chain. Note that such a configuration is not necessary unique (refer to Figure 5 for an example). Level-j chains (for $j > 2$) are defined similarly.

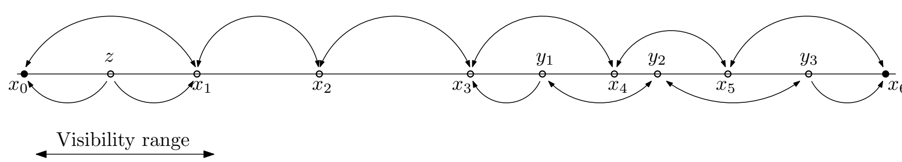

Fig. 5. An example of a primary chain $\{x_0, x_1, \ldots x_6\}$ with two level-2 chains: $\{z\}$ (anchored at x_0 and x_1) and $\{y_1, y_2, y_3\}$ (anchored at x_3 and x_6).

Definition 4 (Secondary Chains and Level-j Chains).
- *The primary chain $\mathcal{C}_1$ is called* level-1 chain*.*
- *A configuration of robots C is a* secondary chain *if it is a mutual chain anchored*

at two robots $x, x' \in \mathcal{C}_1$, and at least one of x and x' is non-faulty. We say that a secondary chain is a level-2 chain.
- *A configuration of robots C is a* level-j chain *if it is a mutual chain anchored at two robots x and x' satisfying the following: there exists an index $j' < j$ such that either x is part of a level-j' chain and x' is part of a level-$(j-1)$ chain, or x is part of a level-$(j-1)$ chain and x' is part of a level-j' chain.*

The convergence of the primary chain can be proven by observing that the behaviour of the robots in the primary chain executing our algorithm (Convergence1D) is equivalent to the behavior they would have if they were executing Algorithm Spreading. Once this is established, convergence follows from Theorem 3. The following lemma shows under what conditions Theorem 3 can be applied to a general mutual chain $Y(t) = \{y_1(t), y_2(t), \ldots, y_k(t)\}$. More specifically, suppose that there exists two real numbers y'_0 and y'_{k+1} such that $y_0(t) = l(y_1(t)) \to y'_0(t)$ and $y_{k+1}(t) = r(y_k(t)) \to y'_{k+1}$ as $t \to \infty$. Then, by executing Convergence1D, $Y(t)$ converges towards an equidistant configuration between $y'_0(t)$ and $y'_{k+1}(t)$.

Lemma 6. *Let $Y(t) = \{y_1(t), y_2(t), \ldots, y_k(t)\}$ be a mutual chain at a size-stable time t, anchored in $y_0(t) = l(y_1(t))$ and $y_{k+1}(t) = r(y_k(t))$, where $y_0(t) \neq y_1(t)$ and $y_{k+1}(t) \neq y_k(t)$. Suppose that there exist two numbers y'_0 and y'_{k+1}, such that $y_0(t) \to y'_0$ and $y_{k+1}(t) \to y'_{k+1}$ as $t \to \infty$. We have that, for all $0 \leq i \leq k+1$,*

$$y_i(t) \to y'_0 + \frac{|y'_{k+1} - y'_0|}{k+1} i \quad \text{as} \quad t \to \infty.$$

Therefore, as $t \to \infty$, the robots in $\{y_1(t), y_2(t), \ldots, y_k(t)\}$ converge to a configuration where the distance between any two consecutive robots is $\frac{|y'_{k+1} - y'_0|}{k+1}$.

Proof. Let $Z(t) = \{z_0(t) = y_0(t), z_1(t), z_2(t), \ldots, z_m(t) = y_{k+1}(t)\}$ be the global configuration of robots at time t, restricted to the interval $[y_0(t), y_{k+1}(t)]$.

By Theorem 4, $Y(t)$ satisfies the following property: for all $1 \leq i \leq k$ and for all $t' \geq t$, $l(y_i(t')) = l(y_i(t))$ and $r(y_i(t')) = r(y_i(t))$. Therefore, even if there is a robot $z_j(t) \in N(y_i(t)) \setminus Y(t)$, the presence of $z_j(t)$ has no impact on the position of $y_i(t+1)$. Consequently, the positions of the robots in $Y(t+1)$, after executing Algorithm Convergence1D on $Y(t)$, are uniquely determined by the positions of the robots in $Y(t)$. Hence, executing Algorithm Convergence1D on $Y(t)$ produces the same result as executing Algorithm Spreading on $Y(t)$, and thus the lemma follows from Theorem 3. □

We now show that the primary chain $\mathcal{C}_1 = \{x'_0, x'_1, x'_2, ..., x'_k\} \subseteq X$, where $x'_0 = x_0$ and $x'_k = x_n$, converges towards a configuration of equidistant robots delimited by its anchors x_0 and x_n.

Theorem 6 (Convergence of the Primary Chain). *Let $\mathcal{C}_1 = \{x'_0, x'_1, x'_2, ..., x'_k\}$ be the primary chain. We have that $x'_0 = x_0$, $x'_k = x_n$ and for all $0 \leq i \leq k$,*

$$x'_i(t) \to \frac{|x_n - x_0|}{k} i \quad \text{as} \quad t \to \infty.$$

Proof. Since $\mathcal{C}_1$ is a mutual chain, the configuration $\{x'_1, x'_2, ..., x'_{k-1}\}$ is also a mutual chain. It is anchored at x'_0 and x'_k, where $x'_0 \neq x'_1$ and $x'_k \neq x'_{k-1}$. Since the anchors $x'_0 = x_0 = 0$ and $x'_k = x_n$ are faulty, they do not move, and the theorem follows directly from Lemma 6. □

We now show that every level-j chain converges towards a configuration of equidistant robots.

Theorem 7 (Convergence of Level-j Chains). *Let $C_j = \{y_1, y_2, \ldots, y_k\}$ be a level-j chain, where $j \geq 1$ is an integer. Let t be a size-stable time. Let $y_0(t) = l(y_1(t))$ and $y_{k+1}(t) = r(y_k(t))$. There exist real numbers y'_0 and y'_{k+1} such that $y_0(t) \to y'_0$ and $y_{k+1}(t) \to y'_{k+1}$ as $t \to \infty$. Moreover, for all $0 \leq i \leq k+1$,*

$$y_i(t) \to y'_0 + \frac{|y'_{k+1} - y'_0|}{k+1} i \quad as \quad t \to \infty.$$

Proof. We proceed by induction on j. By Theorem 6, our statement is true for $j = 1$. Suppose that the theorem is true for all integers from 1 to $j-1$. Consider a level-j chain $C_j = \{y_1, y_2, \ldots, y_k\}$ anchored at $y_0(t) = l(y_1(t))$ and $y_{k+1}(t) = r(y_k(t))$, where t is a size-stable time. By Defintion 4, there exists an index $j' < j$ such that one of the following two statements is true:
- y_0 is part of a level-j' chain and y_{k+1} is part of a level-$(j-1)$ chain, or
- y_0 is part of a level-$(j-1)$ chain and y_{k+1} is part of a level-j' chain.
Without loss of generality, suppose that y_0 is part of a level-j' chain and y_{k+1} is part of a level-$(j-1)$ chain. By the induction hypothesis, there exist two real numbers y'_0 and y'_{k+1} such that $y_0(t) \to y'_0$ and $y_{k+1}(t) \to y_{k+1}$ as $t \to \infty$. The theorem follows from Lemma 6. □

The following lemma states that every robot belongs to some level-j chain. To simplify the presentation, we assume that the faulty robot x_0 is part of the *level*-0 chain $\{x_0\}$ and that the faulty robot x_n is part of the *level*-0 chain $\{x_n\}$.

Lemma 7. *For all size-stable time t and all $0 \leq i \leq n$, $x_i(t) \in X(t)$ belongs to a level-j chain.*

Proof. Suppose that the statement is false. Let $y_1(t)$ be the leftmost robot that does not satisfy the statement. We will derive a contradiction. Since the leftmost robot x_0 is faulty, $l(y_1(t))$ belongs to a mutual chain, say $C(t) = \{x''_1, x''_2, ..., x''_m\}$, where $l(y_1(t)) = x''_\alpha$ for some index $1 \leq \alpha \leq m$. Let $Y = \{y_1, y_2, ..., y_k\}$ be the configuration of robots such that (refer to Figure 6): 1) $y_i(t) = r(y_{i-1}(t))$ for all $2 \leq i \leq k$, 2) $r(y_k(t))$ belongs to a mutual chain, and 3) for all $1 \leq i \leq k$, $y_i(t)$ does not belong to a mutual chain. Observe that the definition of Y allows k to be equal to 1 (in such a case, only items 2) and 3) apply). By construction and by definition of $y_1(t)$,$\{y_1(t), y_2(t), ..., y_k(t)\}$ is not a mutual chain. Therefore, for the rest of the proof, $k \geq 2$. Let $\{z_1, z_2, ..., z_k\}$ be the configuration of robots such that $z_k = y_k$ and $z_i(t) = l(z_{i+1}(t))$ for all $1 \leq i \leq k-1$. Using the same arguments as in the proof of Theorem 5, we get that $x''_\alpha \leq z_1 \leq y_1$ and $y_{i-1} < z_i \leq y_i$ for all $2 \leq i \leq k$. Since $\{y_1(t), y_2(t), ..., y_k(t)\}$ is not a mutual

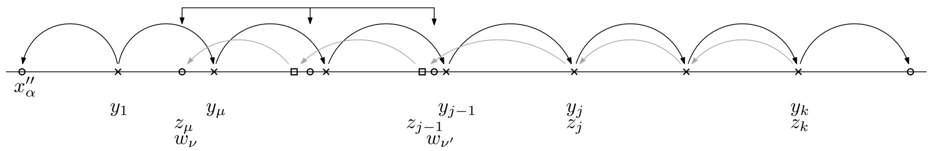

Fig. 6. Illustration of the proof of Lemma 7.

chain, there is an index i such that $z_i(t) \neq y_i(t)$. Let j be the smallest index such that $z_j = y_j$ and $z_{j-1} \neq y_{j-1}$. Suppose there is an index $\gamma < j-1$ such that $z_\gamma(t) = y_\gamma(t)$. Therefore, by the definition of j, $z_i = y_i$ for all $1 \leq i \leq \gamma$. Moreover, x''_α and $r(y_k)$ are part of mutual chains. Therefore, by Theorems 6 and 7, $x''_\alpha(t)$ and $r(y_k)(t)$ converge to a fixed location as $t \to \infty$. Consequently, we get the same contradiction as in the proof of Theorem 5. Hence, for the rest of the proof, assume that $z_i(t) \neq y_i(t)$ for all $1 \leq i < j-1$.

We have the following property (whose proof can be found in [10]).

Property 1. *If, for all $2 \leq i \leq j-1$, $z_i(t)$ does not belong to any mutual chain, then $z_1(t) = l(z_2(t))$ belongs to a mutual chain.*

Consequently, there is an index $1 \leq i \leq j-1$ such that z_i belongs to a mutual chain. Let $1 \leq \mu \leq j-1$ be the largest index such that z_μ belongs to a mutual chain, say $W = \{w_1, w_2, ..., w_{m'}\}$. Let $1 \leq \nu \leq m'$ be the index such that $w_\nu = z_\mu$. We then have another property.

Property 2. $z_{\mu+1} < w_{\nu+1} < y_{\mu+1}$.
Proof: Observe that $w_{\nu+1} = r(w_\nu)$. We must have that $w_{\nu+1} \leq y_{\mu+1}$ and $w_{\nu+1} \geq z_{\mu+1}$, otherwise there would be a contradiction with the fact that $y_{\mu+1} = r(y_\mu)$ and $z_\mu = l(z_{\mu+1})$, respectively. Moreover, by definition, we know that $w_{\nu+1} \neq y_{\mu+1}$ and $w_{\nu+1} \neq z_{\mu+1}$.

By repeating the argument for proving Property 2, we reach the index ν' such that $z_{j-1} < w'_{\nu'} < y_{j-1}$. Observe that $w_{\nu'+1} = r(w_{\nu'}) \leq y_j$ and $w_{\nu'+1} \geq y_j = z_j$, otherwise there would be a contradiction with the facts that $y_j = r(y_{j-1})$ and $z_{j-1} = l(z_j)$, respectively. However, by the definition of Y, y_j is not part of a mutual chain. We get a contradiction. □

The following theorem follows directly from Theorems 6 and 7, and Lemma 7.

Theorem 8 (Global Convergence). *For all $0 \leq i \leq n$, $|x_i(t+1) - x_i(t)| \to 0$ as $t \to \infty$. Therefore, $X(t)$ converges towards a fixed configuration $C^* = \{x_0^*, x_1^*, ..., x_n^*\}$ as $t \to \infty$. The configuration C^* contains a primary chain $\mathcal{C}_1$ anchored at x_0 and x_n. Additionally, there is an integer $\kappa \geq 1$ such that for all $0 \leq i \leq n$, x_i^* belongs to a level-j chain, for some $1 \leq j \leq \kappa$. Moreover, every level-j chain in C^* is a mutual chain of equidistant robots.*

References

1. N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM Journal on Computing*, 36(1): 56–82, 2006.

2. H. Ando and Y. Oasa and I. Suzuki and M. Yamashita. A distributed memoryless point convergence algorithm for mobile robots with limited visibility, *IEEE Trans. Robotics and Automation*,15:5, 818–828, 1999.
3. C. Auger, Z. Bouzid, P. Courtieu, S. Tixeuil, X. Urbain. Certified impossibility results for byzantine-tolerant mobile robots. 15th Int. Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS), 178-190, 2013.
4. Z. Bouzid, S. Das, S. Tixeuil. Gathering of mobile robots tolerating multiple crash faults. 14th Int. Conf. on Distr. Comp. and Networking (ICDCS), 337-346, 2013.
5. Z. Bouzid, M. Gradinariu, S. Tixeuil. Optimal byzantine-resilient convergence in uni-dimensional robot networks. *Theor. Comput. Sci.* 411(34-36): 3154-3168, 2010.
6. Q. Bramas, S. Tixeuil. Wait-free gathering without chirality. 22nd Int. Coll. Struct. Information and Comm. Complexity (SIROCCO), 313-327, 2015.
7. M. Cieliebak, P. Flocchini, G. Prencipe, and N. Santoro. Distributed computing by mobile robots: Gathering. *SIAM Journal on Computing*, 41(4):829–879, 2012.
8. R. Cohen and D. Peleg. Convergence properties of the gravitational algorithms in asynchronous robots systems. *SIAM Journal on Computing*, 34:1516–1528, 2005.
9. R. Cohen, D. Peleg, Local spreading algorithms for autonomous robot systems. *Theoretical Computer Science*, 399(1-2): 71-82, 2008.
10. J.-L. De Carufel, P. Flocchini. Fault-induced dynamics of oblivious robots in a line. arXiv:1707.03492, 2017.
11. X. Défago, M. Gradinariu, S. Messika, P. Raipin Parvédy. Fault-tolerant and self-stabilizing mobile robots gathering. 20th Int. Symposium on Distributed Computing (DISC), 46-60, 2006. Extended version in arXiv:1602.05546, 2016.
12. Y. Dieudonné, O. Labbani-Igbida, and F. Petit. Circle formation of weak mobile robots. *ACM Trans.on Autonomous and Adaptive Systems*, 3(4):16:1–16:20, 2008.
13. Y. Dieudonné, F. Petit. Scatter of robots. *Par. Proc. Lett.* 19(1): 175-184, 2009.
14. P. Flocchini, G. Prencipe, N. Santoro, G. Viglietta. Distributed computing by mobile robots: uniform circle formation, *Distributed Computing*, to appear, 2017.
15. P. Flocchini, G. Prencipe, and N. Santoro. Self-deployment algorithms for mobile sensors on a ring. *Theoretical Computer Science*, 402(1):67–80, 2008.
16. P. Flocchini, G. Prencipe, and N. Santoro. *Distributed Computing by Oblivious Mobile Robots.* Synthesis Lectures on Distributed Computing Theory, Morgan & Claypool, 2012.
17. N. Fujinaga, Y. Yamauchi, H. Ono, S. Kijima, M. Yamashita. Pattern formation by oblivious asynchronous mobile robots. *SIAM J. Computing* 44(3): 740-785, 2015.
18. T. Izumi, M. Gradinariu Potop-Butucaru, S. Tixeuil. Connectivity-preserving scattering of mobile robots with limited visibility. 12th Int. Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS), 319-331, 2010.
19. T. Izumi, S. Souissi, Y. Katayama, N. Inuzuka, X. Défago, K. Wada, and M. Yamashita. The gathering problem for two oblivious robots with unreliable compasses. *SIAM Journal on Computing*, 41(1):26–46, 2012.
20. M. Yamashita and I. Suzuki. Characterizing geometric patterns formable by oblivious anonymous mobile robots. *Theoretical Comp. Sci.*, 411(26–28):2433–2453, 2010.
21. Y. Yamauchi, T. Uehara, S. Kijima, M. Yamashita. Plane formation by synchronous mobile robots in the three dimensional euclidean space. *Journal of the ACM*, 63(3), 2017.