

Conic Formation in Presence of Faulty Robots

Debasish Pattanayak^{1,2} *, Klaus-Tycho Foerster²,
Partha Sarathi Mandal¹, and Stefan Schmid²

¹ Department of Mathematics, IIT Guwahati, Guwahati, India

² Faculty of Computer Science, University of Vienna, Vienna, Austria

Abstract. Pattern formation is one of the most fundamental problems in distributed computing, which has recently received much attention. In this paper, we initiate the study of distributed pattern formation in situations when some robots can be *faulty*. In particular, we consider the well-established *look-compute-move* model with oblivious, anonymous robots. We first present lower bounds and show that any deterministic algorithm takes at least two rounds to form simple patterns in the presence of faulty robots. We then present distributed algorithms for our problem which match this bound, *for conic sections*: in at most two rounds, robots form lines, circles and parabola tolerating $f = 2, 3$ and 4 faults, respectively. For $f = 5$, the target patterns are parabola, hyperbola and ellipse. We show that the resulting pattern includes the f faulty robots in the pattern of n robots, where $n \geq 2f + 1$, and that $f < n < 2f + 1$ robots cannot form such patterns. We conclude by discussing several relaxations and extensions.

Keywords: Distributed Algorithms · Pattern Formation · Conic Formation · Fault Tolerance · Oblivious Mobile Robots

1 Introduction

Self-organizing systems have fascinated researchers for many decades already. These systems are capable of forming an overall order from an initially disordered configuration, using *local* interactions between its parts only. Self-organization arises in many forms, including physical, chemical, and biological systems, and can be based on various processes, from crystallization, over chemical oscillation, to neural circuits [7]. Due to their decentralized and self-healing properties, self-organizing systems are often very robust.

We, in this paper, consider self-organizing systems in the context of *robotics*. In particular, we are interested in the fundamental question of how most simple robots can self-organize into basic patterns. This *pattern formation* problem has already received much attention in the literature [6, 18, 23, 24].

* Visit to University of Vienna is supported by the Overseas Visiting Doctoral Fellowship, 2018 Award No. ODF/2018/001055 by the Science and Engineering Research Board (SERB), Government of India.

A particularly well-studied and challenging model is the *look-compute-move* model, in which each robot, in each round, first observes its environment and then decides on its next move. In the most basic setting, the robots do not have any persistent memory and hence cannot remember information from previous rounds, and they can also not distinguish between the other robots they see: the robots are *oblivious* and *anonymous*. Furthermore, robots cannot communicate. Over the last years, several research lines investigated when robots can and cannot form different patterns [6, 13, 18, 19, 23, 24, 26, 27]. However, existing work on pattern formation shares the assumption that robots are non-faulty.

This paper initiates the study of distributed pattern formation algorithms for scenarios where some robots can be *faulty*: the faulty robots do not move nor act according to a specified protocol or algorithm. The setting with faulty robots is particularly challenging, as the non-faulty robots cannot directly observe which robots are faulty. However, even *indirect* observations seem challenging: since all robots are oblivious, a robot cannot remember patterns from previous rounds, and hence, has no information which robots moved recently. In fact, a robot *per se* does not even know whether the current pattern it observes is the initial configuration or whether some rounds of movements already occurred. What's more, the ability to self-organize into a specific pattern seems to require some coordination or even consensus, which is notoriously hard in faulty environments.

Contributions. This paper considers the design of distributed algorithms for pattern formation of most simple oblivious and potentially *faulty* robots. Our main result is an algorithmic framework that allows robots to form patterns which include faulty robots, in a decentralized and efficient manner. In particular, we do not require robots to identify faulty robots explicitly or to remember previous configurations, but require knowledge of the exact number of faults.

For f faults, we show how to form conic patterns in just *two rounds*, for at least $2f + 1$ robots. We form conic patterns such as line, circle and parabola for $f = 2, 3$ and 4, respectively. For $f = 5$, the target pattern are parabola, hyperbola and ellipse. We also prove that this is optimal: no deterministic algorithm can solve this problem in just one round or with less than $2f + 1$ robots. We further discuss several relaxations of our model and extensions of our results, e.g., considering initial symmetric configurations, having at most f faulty robots, or the impossibility of forming the pattern corresponding to f faults. We also discuss an extension where the robots form a line (a circle) for $f = 3, 4, 5$ ($f = 4, 5$).

Organization. After discussing related work in §1.1, we first provide a formal model in §2, followed by a study of the special case of $f = 1$ faulty robot in §3 to provide some intuition. We then give tight runtime and cardinality lower bounds in §4, and match them for the remaining conic patterns in §5. In §6, we show the algorithmic framework and prove the correctness of our algorithm. After discussing further model variations in §7, we conclude in §8.

1.1 Related Work

Pattern formation is an active area of research [13, 19, 27], however, to the best of our knowledge, we are the first to consider pattern formation in the presence of faults: a fundamental extension. In general, pattern formation allows for exploring the limits of what oblivious robots can compute. A “weak robot” model was introduced by Flocchini et al. [18], where the objective is to determine the minimum capabilities a set of robots need to achieve certain tasks. In general, the tasks include *Gathering* [8, 9, 15], *Convergence* [10, 11], *Pattern Formation* [6, 19, 23], etc. Gathering is a special case of pattern formation, where the target pattern is a point. Gathering has been achieved for robots with multiplicity detection [9, 17]. Most gathering algorithms use the capability of multiplicity detection to form a unique multiplicity point starting from a scattered configuration. In the absence of this capability, it has been proved that gathering is impossible in the semi-synchronous model without any agreement on the coordinate system [22].

The objective of gathering algorithms is only to gather the non-faulty robots, not to form general patterns. Moreover, for the specific case of gathering, some interesting first fault-tolerance studies exist. Agmon and Peleg [1] solve the gathering problem for a single crash-fault. Gathering has been solved with multiple faults [4, 5] with strong multiplicity detection. Next, gathering has also been addressed for robots with weak multiplicity detection tolerating multiple crash faults [3, 21]. For byzantine faults, Auger et al. [2] show impossibility results, and Defago et al. [14] present a self-stabilizing algorithm for gathering. The gathering algorithms only gather non-faulty robots. Since oblivious robots cannot differentiate a faulty robot from a non-faulty one, all the algorithms can be considered to be non-terminating algorithms. In contrast in our paper, we include the faulty robots in the pattern, and as a result, we achieve termination.

Flocchini et al. [18] characterize the role of common knowledge, like agreement on the coordinate system as a requisite for the pattern formation problem, and Yamashita and Suzuki [25] characterize the patterns formable by oblivious robots. Fujinaga et al. [19] present an algorithm using bipartite matching for asynchronous robots. Yamauchi and Yamashita [27] propose a pattern formation algorithm for robots with limited visibility. Das et al. [13] characterize the sequence of patterns that are formable, starting from an arbitrary configuration. Das et al. [12] further extend the sequence of pattern formation problem for luminous robots (robots with visible external persistent memory). As a special pattern formation problem, uniform circle formation has also been considered in the literature [16]. Formation of a plane starting from a three-dimensional configuration has also been solved for synchronous robots [24, 26]. The authors characterize configurations for which plane formation is possible. The existing pattern formation algorithms consider the sequential movement of robots. Since we consider faulty robots in our paper, all the existing algorithms are not adaptable to our cause. A fault-tolerant algorithm has to consider the simultaneous movement of robots and should satisfy the wait-free property to avoid cyclic dependency.

2 Preliminaries

We follow standard model assumptions, inspired by existing work, e.g., [1, 3, 21].

2.1 Model

Each robot is a dimensionless point robot. The robots are homogeneous: they execute the same deterministic algorithm, are indistinguishable and anonymous (no identifiers), oblivious (no persistent memories), and silent (no communication). The robots do not share a common coordinate system and observe others in their own *local coordinate system*. The robots have unlimited visibility, and the determined locations of other robots are precise.

Each robot follows the *look-compute-move* cycle. A robot obtains a snapshot of relative positions of other robots with respect to its position in the *look* state. Based on the snapshot of other robot positions, it decides a destination in the *compute* state. In the *move* state, it moves to the destination and reaches the destination in the same round. This is known as *rigid* robot movement. The scheduler, which activates the robots, follows a fully-synchronous (*FSYNC*) model, i.e., all the robots look at the same time and finish movement in the same round, i.e., each completion of look-compute-move cycle is one round. We consider that the robots are susceptible to crash-faults, i.e., they stop moving after the crash and never recover. Moreover, the number of f faulty robots is known beforehand, and as such, the robots know which types of pattern to form. In particular, we assume that the following four initial conditions hold:

1. All f faulty robots have already crashed initially.
2. All initial configurations are asymmetric.³
3. All robots occupy distinct positions initially.⁴
4. The faulty robots form a convex polygon.

The last assumption needs the faulty robots to be at corners of a convex polygon; the non-faulty robots can lie at any position. The rationale behind the assumption is that four robots forming a triangle with a robot inside the triangle do not correspond to any conic section in \mathbb{R}^2 , similarly, for five robots. For three or more robots, a collinear configuration is addressed in §7. For two robots, the assumption trivially holds.

2.2 Notations

A configuration $\mathcal{C} = \{p_1, p_2, \dots, p_n\}$ is a set of n points on the plane \mathbb{R}^2 , where $p_i = (x_i, y_i)$ is a tuple representing the x - and y -coordinates of the robots. Since each robot is initially located at distinct points, it then holds that $p_i \neq p_j$ for any pair of i and j such that $i \neq j$. f is the number of faulty robots. We will

³ This assumption allows us to have a unique ordering of the robots [8].

⁴ As any set of non-faulty robots that share a position will always perform the same actions from then on and be indistinguishable from each other.

always uphold this condition in our algorithms, except for the case of $f = 1$, where the target pattern is a point. The target patterns are conic sections that satisfy the second degree general equation

$$a_1x^2 + a_2y^2 + a_3xy + a_4x + a_5y + a_6 = 0 .$$

Depending on the values of a_i for $i \in \{1, 2, \dots, 6\}$, the equation represents line, circle, parabola, hyperbola or ellipse. We say a set of points form a conic pattern when they lie on the same conic section. Now, we say the conic pattern passes through the set of points. We denote \mathcal{P} as the length of the pattern and u as the uniform distance.

3 Problem Statement and Intuition

Objective. Given a set of robots on the plane as defined in the model (§2.1), we want the robots to form a conic pattern corresponding to the number of faults.⁵

Point Formation ($f = 1$). To provide some intuition, we start with the case of $f = 1$. For a single faulty robot, we move all the robots to the center of their smallest enclosing circle. If there is a faulty robot in the center, then point formation is achieved. If the faulty robot is somewhere else, we arrive at a configuration with two robot positions. For a configuration with two robot positions, all robots move to the other robot's position. From Assumption 3, we have all the robots at distinct initial positions. Hence the faulty robot is at a different position from the gathered robots. Moving all gathered robots to the faulty robot's position achieves our objective of point formation, as the faulty robot cannot move.

4 Lower Bound

We saw above that there could be situations where only one round suffices, namely, for the case of exactly $n = 2$ and $f = 1$. However, we can show that for $f \geq 2$, at least two rounds are required. For conic patterns (with $f \in \{2, 3, 4, 5\}$), this bound is tight: we will later provide algorithms that terminate in two rounds.

Theorem 1. *For every $f \geq 2$ and every $n \geq f + 3$ holds: Any deterministic algorithm needs more than one round to make a pattern passing through all f faulty robots.*

Proof. Let φ be a deterministic algorithm that forms the pattern with faulty robots. Suppose φ solves the pattern in one step. Two patterns can have at most f common points⁶. Let $\mathcal{C} = \{p_1, p_2, \dots, p_{f+3}\}$ be an initial configuration with $f + 3$ robots such that no $f + 1$ robots are in the same pattern.

⁵ That is a point for $f = 1$, a line for $f = 2$, a circle for $f = 3$, a parabola for $f = 4$, and an ellipse or parabola or hyperbola for $f = 5$.

⁶ Two parabolas intersect at 4 points, which can be the common points between two parabola patterns.

Without loss of generality, consider two sets of f faulty robots at positions $\{p_1, p_2, \dots, p_f\}$ and $\{p_2, p_3, \dots, p_{f+1}\}$ and the corresponding pattern be \wp' and \wp'' , respectively. The f faulty robots do not move. Let $\wp' = \{p'_1, p'_2, \dots, p'_{f+3}\}$ and $\wp'' = \{p''_1, p''_2, \dots, p''_{f+3}\}$. As φ achieves pattern formation in one round, both \wp' and \wp'' should be final. We have $p_{f+1} \neq p'_{f+1}$, $p_{f+2} \neq p'_{f+2}$ and $p_{f+3} \neq p'_{f+3}$, since all robots in \wp' are in the pattern. Similarly, we also have $p_1 \neq p''_1$, $p_{f+2} \neq p''_{f+2}$ and $p_{f+3} \neq p''_{f+3}$ for \wp'' . Since the robots at $\{p_2, \dots, p_f\}$ did not move to form \wp' and \wp'' , these $f - 1$ points are common between \wp' and \wp'' . Since, $p_{f+1} \neq p'_{f+1}$ and $p_{f+1} = p''_{f+1}$, so p_{f+1} cannot be a common point between \wp' and \wp'' . Out of p_{f+2} and p_{f+3} , at most one can be a common point in the pattern, since there are at most f common points. Since φ is deterministic, the destination for robots at p_{f+2} and p_{f+3} remains the same regardless of the destination pattern being \wp' or \wp'' . If $\{p'_{f+2}, p'_{f+3}\} = \{p''_{f+2}, p''_{f+3}\}$, then \wp' and \wp'' have $f + 1$ common points. This is a contradiction since the patterns are different. Hence no deterministic algorithm can solve the pattern formation problem with faults in one round. The arguments hold analogously for $n \geq f + 3$ robots. \square

Next, we show a lower bound on the number of robots required to solve the pattern formation problem. A configuration with exactly f robots is trivially solvable since the f robots are already in the pattern. Note that for $f \geq 2$, $2f + 1 \geq f + 3$ holds.

Theorem 2. *At least $2f + 1$ robots are required to form a pattern passing through f faulty robots for $f \geq 2$.*

Proof. Consider the number of robots to be $f < n < 2f + 1$. Assume that the configuration of these n robots is such that no $f + 1$ robots are in the same pattern. Let φ be a deterministic algorithm, which decides the destination of the robots given a configuration. Since the robots are oblivious, it is impossible to determine which robots are faulty given a configuration. As we consider patterns from the conic section, a pattern corresponding to f can be uniquely determined through f robots.

Let φ decide the target pattern corresponding to a set of f robots for the given configuration \mathcal{C} . So the other $n - f$ robots have to move to the pattern. Since the algorithm cannot determine which robots are faulty, the adversary can always choose the $0 < n - f$ robots to be faulty. Since $n - f \leq f$, none of the robots move. This leads to a stagnated configuration, and the algorithm does not proceed further.

If the algorithm decides a pattern that passes through less than or equal to f points in the configuration, then we choose faulty robots out of the points which are not on the pattern, and we arrive at a configuration where not all the robots are in the same pattern. Now there are at most $n - f \leq f$ robots on the pattern, which is the same as the previous configuration. \square

5 Detailed Algorithms

In this section, we provide the promised two round algorithms for the different configurations. To this end, we first provide algorithmic preliminaries in §5.1.

5.1 Algorithmic Preliminaries

In our algorithms, we will perform case distinctions according to the following three types of configurations:

Definition 1 (*Terminal Configuration*). *A configuration is a terminal configuration if all the robots form the target pattern corresponding to f faulty robots.*

Definition 2 (*Type I Configuration*). *If exactly $n - f$ robots are in the target pattern corresponding to f faulty robots, then it is a Type I configuration.*

A Type I configuration can be symmetric or asymmetric.

Definition 3 (*Type O Configuration*). *If a configuration is not Terminal or Type I, then it is a Type O configuration.*

Note that an initial asymmetric configuration can be a Type I or Type O or Terminal configuration. In the following, we also distinguish the configurations based on the uniform spacing between the robots. We use the uniform positions of the robots as a differentiating factor between faulty and non-faulty.

Definition 4 (*Uniform Configuration*). *A configuration is a uniform configuration if the distance between all consecutive pairs of robots in the configuration along the pattern is the same.*

Definition 5 (*Quasi-Uniform Configuration*). *If a uniform configuration with m uniform positions is occupied by n robots, where $n \leq m \leq 2n$, then it is a quasi-uniform configuration.*

With the assumption (§2.1) that the initial configuration is asymmetric, we can obtain an ordering of the robots using the algorithm by Chaudhuri et al. [8].

Lemma 1. *An asymmetric configuration is orderable. [8]*

Using Lemma 1, we can thus always obtain an ordering among the robots. We use this ordering to determine the target pattern in case of a Type O configuration. In general, having an ordering allows us to have complete agreement on the coordinate system, i.e., the robots agree on the direction and orientation.

Let \mathcal{O} be an ordering of the robots that maps the set of robots to a set of integers $\{1, 2, \dots, n\}$ such that each robot corresponds to an integer. This is the rank of the robot in the ordering. In case of symmetry, two robots can have different orderings. Note that locally, the ordering is unique for a particular robot, but that from a global perspective, different robots can have different orderings.

We use the ordering to determine a target pattern only in cases where there are multiple potential target patterns. We always choose the target pattern passing through the smaller ranked robot in the ordering.

Since the algorithm takes at most two steps to reach a pattern containing all faulty robots, we denote the *initial*, *transitional* and *final* configurations as \mathcal{C}_0 , \mathcal{C}_1 and \mathcal{C}_2 , respectively.

5.2 Algorithm

We first present a general algorithm with two different strategies for the open pattern and closed pattern. Among the conic patterns, line, parabola and hyperbola are open patterns, while circle and ellipse are closed patterns. The target position of the robots depend on this since open conics have two end points, while closed do not have any. The length of the pattern, \mathcal{P} is determined with respect to the pattern being formed. For line, \mathcal{P} is the distance between the two endpoints. For parabola (hyperbola), \mathcal{P} is the length of the parabola (hyperbola) between the points where the latus rectum⁷ of parabola (hyperbola) intersects the parabola (hyperbola). In case of circle and ellipse, \mathcal{P} is the perimeter. We denote u as the uniform distance at which the points on the target pattern are determined. The length u is computed along the pattern. We differentiate between two types of configurations:

Type O Configuration. In this case, we form the target pattern outside the smallest enclosing circle of the configuration. The target pattern can be uniquely determined using the asymmetry of the configuration, since a Type O configuration can only appear in the initial state. The target pattern size is dependent on the diameter of smallest enclosing circle.

- Compute the smallest enclosing circle of the configuration where O is the center and d is the diameter.
- Let A be the location of the robot with the smallest rank in the ordering. If A is O , then we choose the robot with second smallest rank.
- Find point B such that B lies on \overrightarrow{OA} and $|\overline{OB}| = d$.
- The target pattern corresponding to f passes through B .

Now, we show how we determine the target pattern corresponding to the number of faults (see Fig. 1).

$f = 2$: The target line is perpendicular to \overline{OB} and has its midpoint at B with length d (ref. Fig. 1a).

$f = 3$: The target circle passing through B has radius d and center at O (ref. Fig. 1b).

$f = 4$ **and** $f = 5$: The target parabola has its vertex at B and focus at O . The latus rectum of parabola is perpendicular to \overline{OB} and length of latus rectum is $2d$ (ref. Fig. 1c).

⁷ The latus rectum is the line that passes through the focus of the parabola and parallel to the directrix.

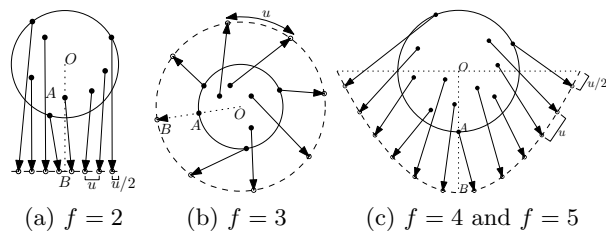


Fig. 1. Target patterns (dashed) for Type O configuration

For an open pattern, we choose the first point at a distance $u/2$ from one end point and place subsequent $n - 1$ points at distance $u = \mathcal{P}/n$ (ref. Fig. 1a, 1c).

Type I Configuration. The target pattern is determined as passing through the robots which are not part of the pattern in the existing configuration. The destinations for robots are at points uniformly positioned at distance u corresponding to the length of the pattern. However, there is a possibility that the existing pattern and target pattern intersect. We hence avoid using the intersection points as target points. If an intersection point is a target point for a value of u , then we choose u' depending on the configuration.

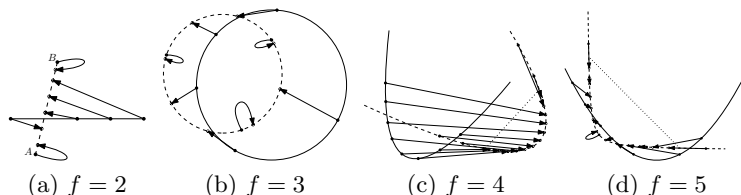


Fig. 2. Target patterns (dashed) for Asymmetric Type I configurations

Asymmetric Type I: The robots are assigned target points according to their rank in the ordering with $u = \mathcal{P}/n$ (ref. Fig. 2). If the target point corresponding to u overlaps with intersection points, then we choose $u' = \mathcal{P}/(n + 1)$. The number of target points with respect to u' is $n + 1$, so we assign two destinations to the robot with highest rank, which can be chosen arbitrarily (ref. Fig. 5). For $f = 5$, the intermediate configuration can be a parabola, hyperbola or an ellipse. Thus, it can also appear as an initial configuration. We show all the transition between ellipse, parabola and hyperbola in the Appendix.

Reflective Symmetric Type I: Let k be the number of robots which lie on the line of symmetry. In this case, we choose, $u = \mathcal{P}/(n + k)$. We get k extra

target points so that we can assign two target points to the robots on the line of symmetry. A robot on one side (say left) of the line of symmetry finds its destination on the same side. Since the robots on the line of symmetry may not have a common left or right, they can choose one of the two symmetric destinations as their target. This also ensures that the next configuration is not completely symmetric (ref. Fig. 3), since only one of the symmetric points would be occupied. On overlap of target pattern points with existing points, we choose u' similar to asymmetric Type I configuration. We refer the reader to the full version of the paper for more details [20].

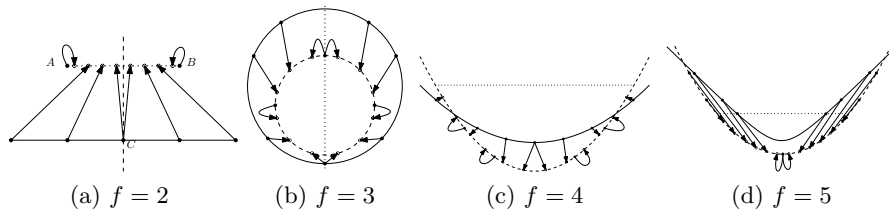


Fig. 3. Target patterns (dashed) for Reflective Symmetric Type I configurations

Rotational Symmetric Type I: This can occur only in the case of lines and circles. In case of a line, the robots move to the side with smaller angle between the target line m and existing pattern l (ref. Fig. 4a). In case of a circle, the configuration is two concentric circles. We assign two destinations for each robot already on the target circle by choosing $u = \mathcal{P}/(n+k)$, where $k = 3$ (ref. Fig. 4b).

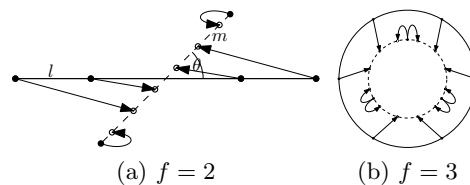


Fig. 4. Target patterns (dashed) for Rotational Symmetric Type I configurations

Choosing Non-overlapping Uniform Points. In case of an asymmetric Type I configuration, we obtain an ordering \mathcal{O} . For $f = 2$, the target line passes through A and B . Let A be the robot with smaller rank among A and B . We choose the set of uniform points at a distance $u/2$ from A towards B . Let D be

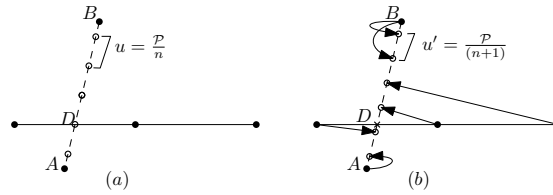


Fig. 5. Choosing Non-overlapping uniform points in the target pattern (dashed)

the intersection point of the existing line in the current configuration and the line through A and B (D is marked as a cross in Fig. 5b). If the uniform points overlap with D (ref. Fig. 5a), then we choose another set of points at uniform distance $u' = \mathcal{P}/(n+1)$. Since, $|\overline{AD}| = 3u/2$, the uniform points overlap with D . We have to choose a set of uniform points corresponding to u' as shown in Fig. 5b. With u' , there are $n+1$ uniform points, the robot B has two potential destinations, one of which can be chosen arbitrarily. We refer the reader to the full version of the paper for the cases of $f = 3, 4, 5$ [20].

6 General Algorithmic Framework

The algorithm broadly has two steps, based on the current robot configuration.

Step 1: Determine the faulty robots.

Step 2: Move to a pattern passing through the faulty robots.

Since the robots are oblivious, they cannot distinguish between Step 1 and Step 2. Hence the properties required for Step 1 have to be applicable to Step 2 and vice versa. For Step 1, a simple process to determine all the faulty robots in a single round is to move all the robots, such that the robots which do not move, will not lie on the pattern points. We determine the pattern points uniformly so that all pattern points are consecutively equidistant along the pattern. This helps us in determining the faulty robots since they would not lie on a pattern point. For Step 2, the pattern determined from the faulty robot positions has to be unique, so that all the robots agree on the pattern. Overall, the algorithm needs to determine a unique pattern such that all robots agree on the pattern, and all robots are required to move to achieve the pattern. We present a transition diagram for configurations in our algorithm in Fig. 6.

Lemma 2. *The destinations of all robots are distinct.*

Proof. We always follow the ordering to determine the destinations for each robot. In the asymmetric cases, we have the ordering, which creates a one-to-one map from the current position of a robot to its destination. In case of symmetric configurations, the robots which are present on the line of symmetry have two potential destinations (from a global perspective) and choose one of them according to their local orientation. Hence the destinations are distinct. \square

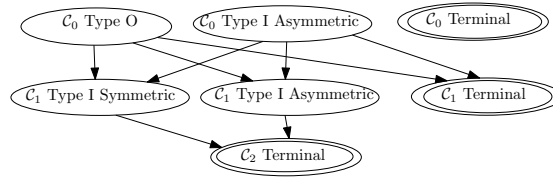


Fig. 6. Transition diagram of configurations

Next, we show that there are no overlapping points between the current configuration and the set of destinations. The set of destinations is the set of all potential destination points for the robots in the current configuration. Since some robots are faulty, two consecutive configurations may have those points as common points. Had the robots been non-faulty, then they would have moved to a point that is not in the current configuration. The proof is in the full version of the paper [20].

Lemma 3. *Given configuration \mathcal{C} and a destination set \mathcal{C}' , we have $\mathcal{C} \cap \mathcal{C}' = \phi$.*

6.1 Determining Faulty Robots and a Pattern

There are two types of initial configurations where we need to determine the faulty robots, i.e., arbitrary configurations and intermediate configurations. The destinations for the robots are such that no point in \mathcal{C}_0 overlaps with any point in \mathcal{C}_1 . For an arbitrary initial configuration, the target pattern is scaled such that no point in the target pattern lies on or inside the smallest enclosing circle of \mathcal{C}_0 . Since \mathcal{C}_0 is asymmetric, we can always uniquely scale the pattern. We can moreover show that a unique pattern exists that passes through all the faulty robots, and obtain the following three lemmas, with proofs of Lemma 4 and 5 in the full version of the paper [20].

Lemma 4. *It takes one round to determine all the faulty robots for a Type O configuration for $f \in \{2, 3, 4, 5\}$.*

Lemma 5. *It takes one round to determine all the faulty robots for an asymmetric Type I configuration for $f \in \{2, 3, 4, 5\}$.*

Lemma 6. *The target pattern passing through the faulty robots in \mathcal{C}_1 can be uniquely determined.*

Proof. The pattern is determined uniquely for a given value of f . For $f = 2$ and 3, the line and circle passing through the points are unique. For $f = 4$, there can be two conjugate parabolas passing through four points. In this case, the parabola with the larger latus rectum is chosen as the target pattern. For $f = 5$, the target pattern is uniquely determined by the five points to be a parabola, hyperbola or an ellipse. Since we assume the faulty robots to form a convex polygon, they can only occupy positions on one side of the hyperbola. From Lemma 3, we know

that the destination points do not have a common point with the points in the previous configuration. From the quasi-uniform configuration, we can determine the robots which are not present at a uniform point. Hence we can determine the faulty robots in \mathcal{C}_1 and the corresponding pattern. \square

6.2 Termination

We can now show that the algorithm terminates, and we can determine the faulty robots in the terminal configuration. Combining Lemma 4, 5 and 6 yields:

Theorem 3. *Starting from any initial asymmetric configuration, the algorithm terminates in at most two rounds.*

Since the algorithm does not do anything for a terminal configuration, we cannot determine the faulty robots if the initial configuration is a terminal configuration. Moreover, in a terminal configuration, our algorithm designs result in the following distribution of robots on the plane starting from a configuration other than the terminal configuration.

Corollary 1. *Starting from a configuration other than the terminal configuration, the non-faulty robots are at uniform pattern points in the terminal configuration.*

Proof. The destinations are always at uniform points spread over the target pattern. So whenever a non-faulty robot moves, it ends up at a uniform pattern point. Note that the resulting configuration may not be uniform due to the faulty robots. The non-faulty robots occupy the uniform points in a quasi-uniform configuration. \square

From Lemma 3 and Corollary 1, we have the following Corollary.

Corollary 2. *The faulty robots can be determined from a terminal configuration unless it is the initial configuration.*

7 Discussion

In the following, we show how to relax our model and extend the previous results in several directions. In particular, we extend the behavior of the algorithm in the absence of the assumption considered in §2. We show that with small modifications to the algorithm, we can subvert some assumptions.

Knowing the number of faults. We extend the definition of Type I configuration to include configurations where at most f robots are not in the pattern in the current configuration. As we need exactly f robots to determine the target pattern, if $f' < f$ are not in the pattern, we choose a target pattern passing through those f' and the first $f - f'$ robots in the ordering to set the pattern.

Initial configuration with reflective symmetry. For a configuration with a single line of symmetry, we can always follow the strategies described for Type

I symmetric configurations in the algorithms from §5. The robots on the line of symmetry have two destinations on either side of the line of symmetry. According to their local orientation, they choose one of the destinations.

Lower order patterns for higher number of faults. We add special cases if the robots are collinear (resp. co-circular) for $f \in \{3, 4, 5\}$ (resp. $f \in \{4, 5\}$). In this case, the robots form a line (resp. circle). If the initial configuration is this situation, then it is impossible to determine the faulty robots. Hence the configuration in the next step becomes an arbitrary configuration. We thus need three steps to achieve pattern formation instead of two.

8 Conclusion

This paper initiated the study of distributed algorithms for pattern formation with faulty robots. In particular, we presented an algorithmic framework that allows solving many basic formation problems in at most two rounds, which is optimal given the lower bound also presented in this paper. We regard our work as a first step and believe it opens several interesting avenues for future research. In particular, it will be interesting to study pattern formation problems for more advanced robots under failures, as well as randomized algorithms. It will also be interesting to generalize our failure model, e.g., to support transient crash faults and byzantine faults.

References

1. Agmon, N., Peleg, D.: Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM J. Comput.* **36**(1), 56–82 (2006)
2. Auger, C., Bouzid, Z., Courtieu, P., Tixeuil, S., Urbain, X.: Certified impossibility results for byzantine-tolerant mobile robots. In: *Stabilization, Safety, and Security of Distributed Systems - 15th Intl. Symposium, SSS 2013, Osaka, Japan, November 13-16, 2013. Proceedings.* pp. 178–190 (2013)
3. Bhagat, S., Mukhopadhyaya, K.: Fault-tolerant gathering of semi-synchronous robots. In: *Proceedings of the 18th International Conference on Distributed Computing and Networking, Hyderabad, India, January 5-7, 2017.* p. 6 (2017)
4. Bouzid, Z., Das, S., Tixeuil, S.: Gathering of mobile robots tolerating multiple crash faults. In: *IEEE 33rd Intl. Conference on Distributed Computing Systems, ICDCS 2013, 8-11 July, 2013, Philadelphia, Pennsylvania, USA.* pp. 337–346 (2013)
5. Bramas, Q., Tixeuil, S.: Wait-free gathering without chirality. In: *SIROCCO 2015, Montserrat, Spain, July 14-16, 2015, Post-Proceedings.* pp. 313–327 (2015)
6. Bramas, Q., Tixeuil, S.: Brief announcement: Probabilistic asynchronous arbitrary pattern formation. In: *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016.* pp. 443–445 (2016)
7. Camazine, S., Deneubourg, J.L., Franks, N.R., Sneyd, J., Bonabeau, E., Theraula, G.: *Self-organization in biological systems.* Princeton university press (2003)
8. Chaudhuri, S.G., Mukhopadhyaya, K.: Leader election and gathering for asynchronous fat robots without common chirality. *J. Discrete Algorithms* **33**, 171–192 (2015)

9. Cieliebak, M., Flocchini, P., Prencipe, G., Santoro, N.: Distributed computing by mobile robots: Gathering. *SIAM J. Comput.* **41**(4), 829–879 (2012)
10. Cohen, R., Peleg, D.: Robot convergence via center-of-gravity algorithms. In: Proc. SIROCCO. pp. 79–88 (2004)
11. Cohen, R., Peleg, D.: Convergence properties of the gravitational algorithm in asynchronous robot systems. *SIAM J. Comput.* **34**(6), 1516–1528 (Jun 2005)
12. Das, S., Flocchini, P., Prencipe, G., Santoro, N.: Forming sequences of patterns with luminous robots. *IEEE Access* **8**, 90577–90597 (2020)
13. Das, S., Flocchini, P., Santoro, N., Yamashita, M.: Forming sequences of geometric patterns with oblivious mobile robots. *Distributed Computing* **28**(2), 131–145 (2015)
14. Défago, X., Potop-Butucaru, M.G., Clément, J., Messika, S., Parvédy, P.R.: Fault and byzantine tolerant self-stabilizing mobile robots gathering - feasibility study -. *CoRR abs/1602.05546* (2016)
15. Flocchini, P., Prencipe, G., Santoro, N.: Distributed Computing by Oblivious Mobile Robots. *Synthesis Lectures on Distributed Computing Theory*, Morgan & Claypool Publishers (2012)
16. Flocchini, P., Prencipe, G., Santoro, N., Viglietta, G.: Distributed computing by mobile robots: uniform circle formation. *Distributed Comput.* **30**(6), 413–457 (2017)
17. Flocchini, P., Prencipe, G., Santoro, N., Widmayer, P.: Gathering of asynchronous robots with limited visibility. *Theor. Comput. Sci.* **337**(1-3), 147–168 (2005)
18. Flocchini, P., Prencipe, G., Santoro, N., Widmayer, P.: Arbitrary pattern formation by asynchronous, anonymous, oblivious robots. *Theor. Comput. Sci.* **407**(1-3), 412–447 (2008)
19. Fujinaga, N., Yamauchi, Y., Ono, H., Kijima, S., Yamashita, M.: Pattern formation by oblivious asynchronous mobile robots. *SIAM J. Comput.* **44**(3), 740–785 (2015)
20. Pattanayak, D., Foerster, K., Mandal, P.S., Schmid, S.: Conic formation in presence of faulty robots. *CoRR abs/2003.01914* (2020), <https://arxiv.org/abs/2003.01914>
21. Pattanayak, D., Mondal, K., Ramesh, H., Mandal, P.S.: Gathering of mobile robots with weak multiplicity detection in presence of crash-faults. *J. Parallel Distrib. Comput.* **123**, 145–155 (2019)
22. Prencipe, G.: Impossibility of gathering by a set of autonomous mobile robots. *Theor. Comput. Sci.* **384**(2-3), 222–231 (Oct 2007)
23. Suzuki, I., Yamashita, M.: Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM J. Comput.* **28**(4), 1347–1363 (1999)
24. Tomita, Y., Yamauchi, Y., Kijima, S., Yamashita, M.: Plane formation by synchronous mobile robots without chirality. In: 21st International Conference on Principles of Distributed Systems, OPODIS 2017, Lisbon, Portugal, December 18–20, 2017. pp. 13:1–13:17 (2017)
25. Yamashita, M., Suzuki, I.: Characterizing geometric patterns formable by oblivious anonymous mobile robots. *Theor. Comput. Sci.* **411**(26-28), 2433–2453 (2010)
26. Yamauchi, Y., Uehara, T., Kijima, S., Yamashita, M.: Plane formation by synchronous mobile robots in the three-dimensional euclidean space. *J. ACM* **64**(3), 16:1–16:43 (2017)
27. Yamauchi, Y., Yamashita, M.: Pattern formation by mobile robots with limited visibility. In: Structural Information and Communication Complexity - 20th International Colloquium, SIROCCO 2013, Ischia, Italy, July 1-3, 2013, Revised Selected Papers. pp. 201–212 (2013)