



HAL
open science

Unreliable Sensors for Reliable Efficient Robots

Adam Heriban, Sébastien Tixeuil

► **To cite this version:**

Adam Heriban, Sébastien Tixeuil. Unreliable Sensors for Reliable Efficient Robots. [Research Report] Sorbonne University. 2021. hal-03219365v2

HAL Id: hal-03219365

<https://hal.sorbonne-universite.fr/hal-03219365v2>

Submitted on 19 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Unreliable Sensors for Reliable Efficient Robots

Adam Heriban* Sébastien Tixeuil*

*Sorbonne University, CNRS, LIP6, France

Abstract

The vast majority of existing Distributed Computing literature about mobile robotic swarms considers *computability* issues: characterizing the set of system hypotheses that enables problem solvability. By contrast, the focus of this work is to investigate *complexity* issues: obtaining quantitative results about a given problem that admits solutions. Our quantitative measurements rely on a newly developed simulation framework to benchmark pen and paper designs.

First, we consider the maximum traveled distance when gathering robots at a given location, not known beforehand (both in the two robots and in the n robots settings) in the classical *OBLLOT* model, for the FSYNC, SSYNC, and ASYNC schedulers. This particular metric appears relevant as it correlates closely to what would be real world fuel consumption. Then, we introduce the possibility of errors in the vision of robots, and assess the behavior of known rendezvous (*aka* two robots gathering) and leader election protocols when sensors are unreliable. We also introduce two new algorithms, one for fuel efficient convergence, and one for leader election, that operate reliably despite unreliable sensors.

1 Introduction

Since the seminal work of Suzuki and Yamashita [41], much research on cooperative mobile robots was aimed at identifying the minimal assumptions (in terms of synchrony, sensing capabilities, environment, etc.) under which basic problems can be solved. A recent state of the art was recently proposed by Flocchini et al. [25].

Robots are modeled as mathematical points in the 2D Euclidean plane and independently execute their own instance of the same algorithm. In the model we consider, robots are anonymous (*i.e.*, they are indistinguishable from each-other), oblivious (*i.e.*, they have no persistent memory of the past is available), and disoriented (*i.e.*, they do not agree on a common coordinate system). The robots operate in Look-Compute-Move cycles. In each cycle, a robot "Looks" at its surroundings and obtains (in its own coordinate system) a snapshot containing the locations of all robots. Based on this visual information, the robot "Computes" a destination location (still in its own coordinate system), and then "Moves" towards the computed location. Since the robots are identical, they all follow the same deterministic algorithm. The algorithm is oblivious if the computed destination in each cycle depends only on the snapshot obtained in the current cycle (and not on stored previous snapshots). The snapshots obtained by the robots are not consistently oriented in any manner (that is, the robots' local coordinate systems do not share a common direction nor a common chirality¹).

The execution model significantly impacts the ability to solve collaborative tasks. Three different levels of synchronization have been commonly considered. The strongest model is the fully-synchronous (FSYNC) model [41], where each phase of each cycle is performed simultaneously by all robots. The semi-synchronous (SSYNC) model [41] considers that time

¹Chirality denotes the ability to distinguish left from right.

is discretized into rounds, and that in each round an arbitrary yet non-empty subset of the robots are active. The robots that are active in a particular round perform exactly one atomic LOOK-COMPUTE-MOVE cycle in that round. The weakest model is the asynchronous (ASYNC) model [26, 27], which allows arbitrary delays between the LOOK, COMPUTE and MOVE phases, and the movement itself may take an arbitrary amount of time. It is assumed that the scheduler (seen as an adversary) is fair in the sense that in each execution, every robot is activated infinitely often.

1.1 Previous works and Motivations

An important shortcoming of the robot model introduced by Suzuki and Yamashita [41] with respect to real-world implementation of mobile robot algorithms is the assumption that both the vision sensors and the actuation motors are perfect. More specifically, the model assumes that robots have an infinite vision range, and can sense the position of other robots relatively to theirs with infinite accuracy. Robots are also usually able to reach their target with infinite movement precision (with respect to the angle to the target).

Several attempts have been made to make the *OBLLOT* model more realistic, *e.g.* by limiting the range of sensors through the limited visibility model [1, 28, 29], by allowing the sensors to miss other robots [31], by using inaccurate sensors [13, 28, 29, 35, 43], or by discarding the hypothesis that robots are transparent [32, 34].

However, many attempts are hindered by increased complexity due to manually proving algorithms in those more complex settings. For instance, to our knowledge, the consequences of error-prone vision have only been studied through very simple problems: **Gathering** and **Convergence** [13, 28, 29, 35, 43].

To allow more complex problems to be studied considering more realistic settings, it appears necessary to favor a machine-helped approach.

Formal methods encompass a long-lasting path of research that is meant to overcome errors of human origin. Unsurprisingly, this mechanized approach to protocol correctness was used in the context of mobile robots [3, 4, 5, 6, 7, 8, 9, 14, 19, 20, 37, 38, 39, 40].

When robots move freely in a continuous two-dimensional Euclidean space, to the best of our knowledge, the only formal framework available is Pactole². Pactole relies on higher-order logic to certify impossibility results [3, 6, 14], as well as the correctness of algorithms [5, 15] in the FSYNC and SSYNC models, possibly for an arbitrary number of robots (hence in a scalable manner). Pactole was recently extended by Balabonski *et al.* [4] to handle the ASYNC model, thanks to its modular design. However, in its current form, Pactole lacks automation; that is, in order to prove a result formally, one still has to write the proof (that is automatically verified), which requires expertise both in Coq (the language Pactole is based upon) and about the mathematical and logical arguments one should use to complete the proof.

On the other hand, model checking and its derivatives (automatic program synthesis, parameterized model checking) hint at more automation once a suitable model has been defined with the input language of the model checker. In particular, model-checking proved useful to find bugs (usually in the ASYNC setting) [8, 22, 23] and to formally check the correctness of published algorithms [8, 19, 20, 38]. Automatic program synthesis [9, 37] was used to obtain automatically algorithms that are "correct-by-design". However, those approaches are limited to instances with few robots. Generalizing them to an arbitrary number of robots with similar models is doubtful as Sangnier *et al.* [40] proved that safety and reachability problems are undecidable in the parameterized case with default models. Another limitation of the above approaches is that they *only* consider cases where mobile robots *evolve in a discrete space* (*i.e.*, graph). This limitation is due to the model used, that closely matches the original execution

²<http://pactole.lri.fr>

model by Suzuki and Yamashita [41]. As a computer can only model a finite set of locations, a continuous 2D Euclidean space cannot be expressed in this model.

Overall, the only way to obtain automated proofs of correctness in the continuous space context through model checking is to use a more abstract model [18, 19], which commands writing additional handwritten theorems to assess its relevance in the original model. Overall, using formal methods for complex algorithms in realistic settings requires a substantial effort that may be out of reach when one simply wants to assess the feasibility of an algorithmic design.

Furthermore, these approaches currently only address whether the added constraints enable the construction of counter-examples for a given task, and, to the best of our knowledge, do not address the important issue of performance degradation, or, in the cases where counter-examples do appear, the likelihood of their appearance and their impact.

In fact, an overwhelming majority of the research on mobile robotic swarms has focused on proving, under a given set of conditions, whether there exists a counter example to a given solution proposal for a problem. On the other hand, the practical efficiency of a given algorithm (with respect to real-world criteria such as fuel consumption) was rarely studied by the Distributed Computing community, albeit being of paramount importance to the Robotics community [2, 45]. Fuel-constrained robots have been considered in the discrete graph context, for both exploration [24] and distributed package delivery [11], but, to our knowledge, no study considered the two-dimensional Euclidean space model that was promoted by Suzuki and Yamashita [41]. A possible explanation for this situation is that the more complex the algorithm (or the system settings), the more difficult it becomes to rigorously find the worst possible executions.

We investigate another approach: since our goal is to bridge the gap between theoretical mobile robots, and actual robotics, we move one step towards robotics and use a very common tool: simulation. First, robot simulators, such as Gazebo³, are industry standard tools for designing physical robots. Then, simulating mobile robots is not a new idea, and has been tried since the very beginning of mobile robots [1].

Our goal is to design and implement a practical simulator for networks of mobile robots that is focused on finding counter-examples and monitoring network behavior, rather than proving algorithms or providing a visual representation. Our vision is that this tool is especially useful in the early stages of algorithm design to eliminate obviously wrong paths, and detect anomalies. It should not be seen as a replacement for formal tools, but as a replacement for researcher intuition when working on a mobile robot network model or algorithm.

As such, the simulator should be easy to use, understand and modify by any Distributed Computing researcher in order to include any new algorithm or model. It should also be capable of monitoring network behavior and output quantitative data points to assess real world performance, according to a given set of metrics, as well as enabling comparison with previously proven algorithms in a given setting.

We first focus on the known limitations of this approach and highlight the difficulty of encoding victory and defeat conditions for the computed executions, and how it impacts our ability to reliably detect counter-examples, as well as the expected consequences of working in a discretized Euclidean space, such as the impossibility to distinguish **Convergence** and **Gathering**.

1.2 Our Contribution

In this paper, we design and implement a practical simulator for mobile robotic swarms evolving in a two-dimensional Euclidean space. To circumvent the obvious problem of an infinite number of initial positions, our simulation framework is based on the Monte Carlo method for choosing initial configurations [36].

³<http://gazebosim.org/>

We first benchmark our simulation framework using a well known problem in the domain: rendezvous. Rendezvous mandates that two robots gather in finite time at the same location, not known beforehand. There exists a number of rendezvous solution for various settings, yet our simulation framework enables fair quantitative comparison. We choose the fuel metric (*a.k.a.* total traveled distance) under various system conditions: FSYNC, SSYNC, and ASYNC schedulers with or without rigid motion.

We then assess the impact of inaccurate visibility sensors on two milestone algorithms: the Center-of-Gravity convergence algorithm [41] for two robots, and the Geoleader election algorithm [10]. It turns out that their behavior is significantly impacted by even small inaccuracies.

To address the shortcoming identified by our simulations in the literature, we design a new two-color, fuel-efficient, convergence algorithm for the ASYNC scheduler, and an improved leader election algorithm that is resilient to inaccurate vision. Both proposal are similarly benchmarked with our simulation framework.

The rest of the paper is organized as follows. Section 2 presents the core technicalities underlying our simulation framework, and its limitations through the problems of *OBLLOT* FSYNC **Convergence**, and **Geoleader Election**. Section 3 demonstrates how the framework can be used for the purpose of performance evaluation, while Section 4 show how realistic error models can be integrated into the entire evaluation process. Section 5 introduces two new algorithms, one for fuel efficient convergence, and one for leader election with unreliable sensors. Finally, Section 6 provides concluding remarks.

2 Monte-Carlo Simulation of Mobile Robots

2.1 Overview of the Framework

Our simulation framework is written from scratch using Python 3, ensuring a large compatibility across executing platforms. Our design goal is to remain as close as possible to the theoretical model of Suzuki and Yamashita [41], in order to maximize readability and usability by the mobile robot distributed computing community.

Each mobile entity is thus encapsulated as an instance of the Robot class. In the case of the basic *OBLLOT* model, robots have the following properties:

- A unique name.
- x and y coordinates in the Euclidean plane.
- A snapshot list of Robots that contains visible Robots.
- A target, which is a 2-tuple of the x and y coordinates of the target destination.

The Robot class also provides three methods:

- The LOOK method uses the network as an input. It creates a list of the visible Robots in the network and assigns it to snapshot.
- The COMPUTE method uses snapshot to compute and assign target, according to the algorithm we want to evaluate.
- The MOVE method updates x and y according to target.

This is summarized in figure 1

Because robots are anonymous, name cannot be used for computing purposes, and is simply a way for the scheduler to reliably monitor the robots in the network. Similarly, robots cannot

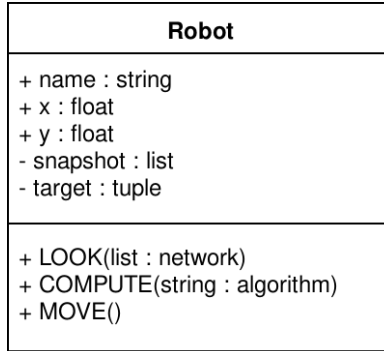


Figure 1 – Robot Class

use x and y directly as they are disoriented. The simulation consists of two parts: an initializing sequence and a loop. The *initializing sequence* creates a *network* list, which contains all robots, according to simulation parameters. To circumvent the problem of the infinite number of possible initial positions, our simulation framework is based on the Monte-Carlo method for choosing initial configurations [36]. So, unless otherwise specified, the initial location of each robot is chosen uniformly at random within the bounds of the type used to represent positions. Using the Monte-Carlo method allows us to both minimize biases in the initial parameters, and arbitrarily increase the precision of the simulation by simply increasing the number of simulations. For each iteration of the *main loop*, a scheduling function is executed once. In the case of FSYNC, for each loop iteration, all robots in the network simultaneously perform a LOOK, then simultaneously perform a COMPUTE, and then simultaneously perform a MOVE. Using different schedulers, such as SSYNC or ASYNC only requires changing the scheduling function: SSYNC creates a non empty list of robots to be activated for a whole cycle, and ASYNC picks a single robot to be activated for a single phase. The loop terminates whenever a *victory* condition holds, which confirms the algorithm completed its intended task. In the case where an algorithm may fail, a *defeat* condition can also be used. For practical reasons, the loop has a maximum number of iterations. However, reaching this maximum should not be interpreted as either a failure or a success.

2.2 Scheduling

Modeling the FSYNC scheduler can be trivially done by performing all LOOK operations, then all COMPUTE operations, then all MOVE operations. For the ASYNC and SSYNC schedulers, we rely on randomness to test as many executions as possible. To model the SSYNC scheduler, for each time step, we chose a non-empty subset of the network uniformly at random and perform a full cycle.

To model the ASYNC scheduler, we chose one robot uniformly at random and perform its next operation⁴. In the case of the ASYNC scheduler, we must also consider what happens if a robot performs a LOOK operation while another robot is moving. the *OBLLOT* model usually considers that an adversary can chose the perceived location of the second robot to be anywhere between its initial position and its destination (on a straight line). Modeling this behavior could be easily done by changing the perceived coordinates in the LOOK operation uniformly at random between the location and the target of the perceived robot (on a straight line). However, existing literature about the ASYNC model shows that the most problematic scenarios appear when the outdated position perceived for a robot is its initial location. With our simulation framework, we also observed that always choosing the initial location when

⁴Note that this model does not explicitly include simultaneous operations: we consider that the output of two simultaneous events E_1 and E_2 can be either the output of E_1 then E_2 or the output of E_2 then E_1 .

observing a given robot while it is in its MOVE phase yielded the most adversarial results, so, while our framework is able to simulate both perceptions, we assume this adversarial behavior in the sequel.

For all schedulers, our simulation framework supports both the rigid and the non-rigid settings. The rigid setting mandates a robot that selected a distinct target in the COMPUTE phase to always reach it in the MOVE phase. The non-rigid setting partially removes this condition: the robot may be stopped by the scheduler before it reaches the target, but not before it traverses a distance of at least δ , for some $\delta > 0$.

2.3 Simulation Conditions

Our framework uses Monte-Carlo simulation for both the initial conditions and the scheduling. This means we can perform an arbitrarily large number of simulations, which in turn induces an arbitrarily more precise simulation. Therefore any criterion on either time, number of iterations, or precision is equivalent. Unless specified otherwise, 4 simulation threads are run in parallel, for one hour, on a modern quad-core CPU, after which results are merged and analyzed. We use the PyPy3 JIT compiler instead of the CPython interpreter, for better performance. Results of the 4 simulations are then compiled and analyzed.

2.4 Comparison with Existing Simulators

We found two noteworthy simulators for mobile robots: Sycamore and JBotSim.

Sycamore is a Java program focused explicitly on mobile robots. However, it appears to be far more complex to build, use and modify than our proposal. Moreover, the latest version we could find seems to date back from 2016, and requires versions of Java that are no longer supported.

*JBotSim*⁵ is a Java library for simulating distributed networks in general. While it appears to be able to simulate *OBLLOT* robots, it is not designed to do so. So, one has to dig into the intricacies of the simulator to emulate basic mobile robot settings.

We also found a third Java-based simulator, named *oblot-sim*⁶. We are, however, unsure of its provenance and design goals.

All three simulators emphasize real-time visualization of executed algorithms through a complete graphical interface. Our proposal focuses on extremely simple quantitative simulation. In its current version, a complete instance of the simulator requires only five separate files for a total of less than 30KB of code (The sources for JBotSim and Sycamore weigh 3MB and 4.8MB, respectively). We also believe that using Python instead of Java greatly improves portability and ease of understanding, which in turns allows researchers to more easily implement and test unusual settings.

In short, our goal is not to visualize executions in real-time, but to simulate as many executions as possible to process their outcome.

2.5 Limitations of the Simulation

While the initial approach described in previous sections may seem sound and simple enough to work with, it results in two distinct problems. As stated previously, our objective with robot simulation is to reliably provide counter-examples whenever they may occur. This requires reliably detecting problematic executions, which is difficult for two reasons. First, success and defeat conditions for most mobile robot algorithms are written in a way that might not be directly usable in a computer simulation. Then, we show that issues predictably arise due to the nature

⁵<https://jbotsim.io>

⁶<https://github.com/werner291/oblot-sim>

of discretized floating point numbers compared to "true" real numbers used in mathematical models.

Halting the Simulation: Victory and Defeat Conditions: One of the goals of our simulation framework is to find counter-examples for a given algorithm and setting. To do so, we need to simulate the evolution of the network until one of two things happen:

- A sufficient condition has been met. This implies that the current execution is successful, and a new simulation with a different initial configuration should begin. This is called a *victory condition*.
- A necessary condition has been violated. This implies that the current execution constitutes a counter-example. This is called a *defeat condition*.

We illustrate the difficulty of using and defining such conditions in practice through the example of one of the most fundamental problems in the context of mobile robots: **Gathering**.

The common victory condition for **Gathering** is the following, for two robots r_1 and r_2 :

Condition 2.1 (Theoretical **Gathering** Victory). ***Gathering** is achieved if and only if, for any pair of robots in the network, the distance between the two robots is eventually always zero. This can also be written more formally as $\exists t_0 \in \mathbb{R}_{\geq 0} : \forall t_1 \geq t_0, \forall (r_1, r_2) | r_1 r_2 |_{t_1} = 0$*

In the previous condition, $|r_1 r_2|_t$ denotes the distance between r_1 and r_2 at time t in the current execution.

However, this particular condition would require the ability for the simulator to infinitely simulate the future of the network, which is obviously impossible. Moreover, the matching defeat condition is unusable for similar reasons:

$$\nexists t_0 \in \mathbb{R}_{\geq 0} : \forall t_1 \geq t_0, \forall (r_1, r_2) | r_1 r_2 |_{t_1} = 0$$

or

$$\forall t_0 \in \mathbb{R}_{\geq 0} : \exists t_1 \geq t_0, \exists (r_1, r_2) | r_1 r_2 |_{t_1} \neq 0$$

We instead define a more practical defeat condition:

Condition 2.2 (Practical **Gathering** Defeat). $\exists (t_0, t_1) \in (\mathbb{R}_{\geq 0})^2 : t_1 > t_0, inputs(t_0) = inputs(t_1), \exists t \in [t_0, t_1] / \exists (r_1, r_2) | r_1 r_2 |_t \neq 0$

Where $inputs(t)$ is the set of all input parameters relevant to the algorithm. This is different from the configuration, which would contain *all* parameters of the network at a given point of the execution.

This input set is used as a practical way to detect cycles in the execution. For a deterministic algorithm, if all inputs of the algorithm are identical to a previously encountered set of inputs, then a cycle has been found.

The input set we use must be chosen such that for two sets S_1 and S_2 , $S_1(t) = S_2(t) \implies \forall S_1(t+1), \exists S_2(t+1) : S_1(t+1) = S_2(t+1)$. In other words, regardless of the scheduling, two identical sets should not be able to generate different sets.

Theorem 2.1. *For two robots executing a deterministic algorithm, if condition 2.2 is true then condition 2.1 is false.*

Proof. For a deterministic algorithm, if condition 2.2 is true, there exists a scheduling starting from the initial configuration that reaches $inputs(t_0)$ and $inputs(t_1)$. Because $inputs(t_0) = inputs(t_1)$, there exists a cycle containing non-gathered configurations. Then the adversary scheduler can repeat this cycle infinitely, and condition 2.1 is false. \square

Theorem 2.2. *If the number of input sets is finite, then for two robots executing a deterministic algorithm, if condition 2.1 is false, then condition 2.2 is true.*

Proof. Any scheduling is infinite. So, if the total number of input sets is finite, then every scheduling contains at least one cycle. If condition 2.1 is false, then there are no non-gathered cycles, so there is at least one gathered cycle that must be repeated, and condition 2.2 is true. \square

One may naively want to use a similar reasoning to define a sufficient victory condition:

Condition 2.3 (Naive **Gathering** Victory). $\exists(t_0, t_1) \in \mathbb{R}_{\geq 0}^2 : t_1 > t_0, \text{inputs}(t_0) = \text{inputs}(t_1), \forall t \in [t_0, t_1], \forall(r_1, r_2) |r_1 r_2|_t = 0$

However, this condition ignores the fact that the scheduler may be able to not repeat this cycle by carefully choosing the activation order of the robots. A proper condition that is usable regardless of the scheduler is the following:

Condition 2.4 (Practical **Gathering** Victory). $\forall(r_1, r_2) \exists t_0 \in \mathbb{R}_{\geq 0} : |r_1 r_2|_{t_0} = 0 \wedge \forall \mathcal{S}, \exists t_1 > t_0 : \text{inputs}(t_0) = \text{inputs}(t_1), \forall t \in [t_0, t_1], |r_1 r_2|_t = 0$

With \mathcal{S} a scheduling. In other words, there exists a time after which all robots are stuck in gathered cycles.

Analyzing configurations and finding cycles in the execution is not an issue for our simulator. The main difficulty lies in our ability to properly model the configuration using the input set. If the set is too restrictive and omits relevant parameters, then we find cycles that do not actually exist. Similarly, a set that is not restrictive enough may hide actual cycles. This depends on both the robot model and the algorithm used to solve the problem.

In the case of **Rendezvous** or **Gathering** for two robots, the standard algorithm [41] for the FSYNC scheduler targets the midpoint between the two robots and is described in algorithm 2.1.

Algorithm 2.1 Basic FSYNC Rendezvous

target[0] = (x + snapshot[0].x)/2
target[1] = (y + snapshot[0].y)/2

In the Euclidean space, the number of configurations appears to be infinite. Because robots are disoriented, the algorithm uses no information on distance, or coordinate systems, so that all configurations are identical. Then, the input set is actually empty. This implies that an algorithm succeeds if and only if the network is gathered after the first activation of both robots. Otherwise, the defeat condition is immediately true for rigid movement.

For the sake of providing a second example, let us consider that robots are endowed with weak local multiplicity detection, meaning that they can distinguish a non-gathered configuration from a gathered configuration. This allows us to modify the initial algorithm to algorithm 2.2.

Algorithm 2.2 FSYNC Rendezvous with Multiplicity Detection

if \neg gathered
 target[0] = (x + snapshot[0].x)/2
 target[1] = (y + snapshot[0].y)/2

In this case, the gathered state is a relevant input parameter, and should be included in the input set. Now, all gathered configurations are considered identical and all non-gathered configurations are considered identical. This means that the robots must still gather after the first activation. However, while this was already considered a cycle with the empty set, if robots are now gathered, the input set is different and no cycle has yet been reached. The first cycle

is reached after the second activation. If the robots remain gathered, then this is a gathered cycle and should not trigger the defeat condition. However, if for some reason the robots were to separate after the second activation, this would constitute a non-gathered cycle with the first input set, and the defeat condition would be triggered.

Using this reasoning, we check our simulator against our two-color ASYNC algorithm [30] and the two-color SSYNC algorithm from Viglietta [42]. For Heriban two-color, we accurately find no counter-example, and all executions lead to the victory condition in ASYNC, SSYNC and FSYNC. For Viglietta two-color, we accurately find no counter-example and all executions lead to the victory condition in SSYNC and FSYNC, and we find counter-examples that trigger the defeat condition in ASYNC.

We perform a similar study for a weaker version of **Gathering**, called **Convergence**. The common condition for **Convergence** is the following:

Condition 2.5 (Theoretical **Convergence** Victory). ***Convergence** is achieved if and only if, for any distance ε greater than zero, the distance between any pair of robots is eventually always smaller than ε .*

This can also be written more formally as $\forall \varepsilon \in \mathbb{R}_{>0}, \exists t_0 \in \mathbb{R}_{\geq 0} : \forall t_1 \geq t_0, \forall (r_1, r_2) | r_1 r_2 |_{t_1} \leq \varepsilon$

Note that, as we expect, **Gathering** implies **Convergence**, but **Convergence** does not imply **Gathering**. In this case, the distance between the two robots is a relevant parameter to check whether or not the problem is solved. However, since it does not change the behavior of the algorithm, it is still not part of the input set.

We define the following defeat condition:

Condition 2.6 (Practical **Convergence** Defeat). $\exists (r_1, r_2) : \exists (t_0, t_1) \in (\mathbb{R}_{\geq 0})^2 : t_1 > t_0 \wedge \text{inputs}(t_0) = \text{inputs}(t_1) \wedge 0 < |r_1 r_2|_{t_0} \leq |r_1 r_2|_{t_1}$

Theorem 2.3. *For a deterministic algorithm, if condition 2.6 is true, then condition 2.5 is false.*

Proof. Similarly to **Gathering**, this condition implies a cycle where distance does not decrease, so the adversary scheduler can repeat it infinitely and prevent **Convergence**. \square

This does *not* imply that the distance between the two robots must always be strictly decreasing in the general case, as this would neither be a sufficient nor a necessary condition. Because ε can be infinitely small, we cannot choose the 'right' ε to properly define a victory condition.

The Consequences of the Discretized Euclidean Plane: While it is tempting to define a victory condition similar to that of **Gathering**, the question of ε remains. Floating point numbers are obviously incapable of infinite precision. So, because any number greater than zero is a valid choice, if ε is smaller than the minimum positive number that can be represented in the chosen floating point precision, it cannot be distinguished from a true zero. This implies that small enough distances between two robots cannot be distinguished from a gathered state. So, it is intrinsically impossible to distinguish **Convergence** from actual **Gathering**.

Let us modify algorithm 2.1 so that both robots move towards the midpoint, but only move a distance of $\frac{|r_1 r_2|}{2} - \frac{\delta}{2}$ instead of $\frac{|r_1 r_2|}{2}$. In theory, this algorithm does not lead to **Rendezvous**, as robots reach a distance of δ after their first activation. However, if δ is small enough, the precision of floating point numbers is such that $\frac{|r_1 r_2|}{2} - \frac{\delta}{2}$ and $\frac{|r_1 r_2|}{2}$ appear identical, and the distance $|r_1 r_2|$ appears to be zero. This is essentially a **Convergence** algorithm that is fast enough to be mistaken for a **Rendezvous** algorithm. In practice, there is very little that can be done against this sort of behavior and conditions for **Gathering** should not be considered reliable.

On the other hand, under different circumstances, the discrete nature of the simulation can instead lead theoretically good executions to fail in practice. Let us consider a network of two robots r_1 and r_2 such that r_2 does not move, and r_1 moves to the midpoint. This should trivially lead to **Convergence**. Let us now assume that $r_1.y = r_2.y$, and that $r_1.x$ and $r_2.x$ are such that $r_2.x$ is the smallest float greater than $r_1.x$. This possibly leads to $\frac{r_1.x + r_2.x}{2} = r_1.x$, so r_1 stops moving and the defeat condition for **Convergence** is wrongly activated.

We test this by setting $r_1.y = r_2.y = 0$, picking $r_1.x$ at random in $[0, 1]$ and picking $r_2.x$ at random in $[2, 3]$ so that $r_1.x < r_2.x$.

In the first case, r_1 moves to the midpoint and r_2 does not move. This results in approximately 37.5% of one million attempts wrongly failing **Convergence**.

In the second case, r_2 moves to the midpoint and r_1 does not move. This results in approximately 25.0% of one million attempts wrongly failing **Convergence**.

This asymmetry may be explained by biases in the binary64 approximation. Regardless, this is a real, hard to predict problem with a non-negligible chance of happening and requires careful analysis of found counter-examples.

Problems with limited float precision also appear when simulating **Geoleader Election**.

Geoleader Election is successful if, given a set of robots, each with their own coordinate system, robots can all deterministically agree on a same robot, called the **GEOLEADER**.

Geoleader Election is known to be impossible in the general case [21] because of possible symmetries in the network. In practice, this impossibility is circumvented using randomized algorithms to break such symmetries. Let us consider the state-of-the-art algorithm 2.3 by Canepa and Gradinariu Potop-Butucaru [10] for three robots.

Algorithm 2.3 Original Leader Election Algorithm by Canepa and Gradinariu Potop-Butucaru [10] for Three Robots

```

Compute the angles between two robots
if my_angle is the smallest
    Become LEADER
    Exit
else if my_angle is not the smallest, but the other two are identical
    Become LEADER
    Exit
else if All angles are identical
    Perform a Bernoulli trial with a probability of winning of  $p = \frac{1}{3}$ 
    if Trial won
        Move perpendicular to the opposite side of the triangle in opposite direction

```

For this particular algorithm, there are three cases:

1. The common case, where one angle is greater than the two others.
2. A rare case where two angles are identical, and the third one is smaller.
3. The rarest case where all angles are identical. In that case, a Bernoulli trial is required to degrade to the other cases.

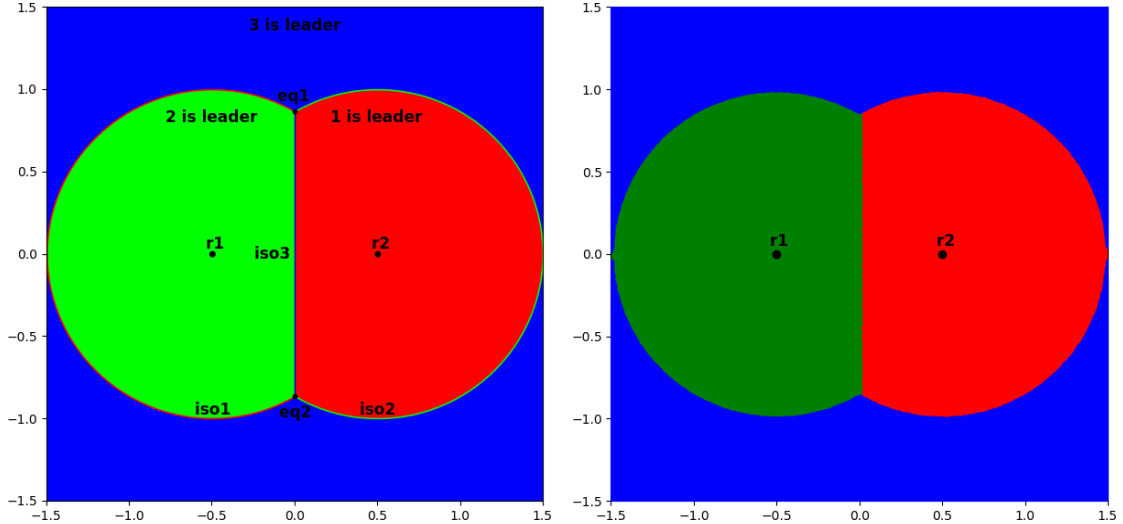
Let us assume a network of three robots, $[r_1, r_2, r_3]$, such that r_1 is placed at coordinates $(-0.5, 0)$, and r_2 at $(0.5, 0)$.

We show where each case appears in figure 2a). The third case occurs if r_3 is at $(0, \pm \frac{\sqrt{3}}{2})$, which are noted as points *eq1* and *eq2*. Positions of r_3 that lead to the second case are noted as *iso1*, *iso2*, and *iso3*.

However, it is *not* possible, using floating point numbers, to have x such that $x^2 = 3$. It is then impossible, regardless of the quality of the simulation, to place r_3 on $eq1$ or $eq2$, despite being possible in theory.

Similarly, an infinitely large number of points mathematically located on the circular arcs of the second case cannot be represented properly using floating point numbers.

To test this, each robot is given a new property 'Leader', which is a string containing the name of the **LEADER** robot. We perform the simulation and display the results in figure 2b.



(a) **LEADER** Depending on the Location of r_3 . Red, green and blue represent r_1 , r_2 and r_3 , respectively.

(b) Simulation for 3-robot **Leader Election** with Perfect Vision Sensors. No isosceles or equilateral point was found.

As we predicted, the fact that real numbers cannot be properly represented in our discrete, floating point space prevents the simulator from finding the known counter-example in the case of 3-robots **Leader Election**. Furthermore, the three circular arcs on which the second case occurs have a combined surface theoretically equal to zero. Therefore, they are statistically impossible to find using our Monte-Carlo simulation.

However, it should be noted that, even in a world of perfect sensors, building an equilateral triangle would require placing the third robot with physically impossible precision. So, while this counter-example exists from a mathematical standpoint, it could never occur in a more realistic setting. So, when considering practical robots, this could be considered a minor issue.

On the contrary, the use of a discretized Euclidean space could be viewed as massive advantage compared to the regular, continuous model, as it makes the inherently unrealistic hypothesis of robots being able to store and process snapshots of infinite precision. In this approximated context, snapshots have a known, maximum size, depending on the chosen precision for the coordinates of other robots.

So, in this context, storing a snapshot for a full cycle becomes a trivial matter, and using algorithm **SyncSim** described by Das *et al.* [16, 17], to simulate an FSYNC scheduling under an ASYNC scheduler, becomes possible without additional unrealistic hypotheses.

As a result, we believe designing algorithms that properly solve problems in the context of a discretized Euclidean space should be a priority, as it would allow mobile robots to only need to function using the FSYNC scheduler, and would remove the unrealistic requirement of infinite precision. One such algorithm is shown in Section 5.

3 Fuel Efficiency in the Usual Settings

The overwhelming majority of the mobile robots research has focused on proving, under a given set of conditions, whether there exists a counter example to a given problem. On the other hand, the practical efficiency of a given algorithm (with respect to real-world criteria such as fuel consumption) was rarely studied by the distributed computing community, albeit commanded by the robotics community [2, 45].

Fuel-constrained robots have been considered in the discrete graph context, for both exploration [24] and distributed package delivery [11]. However, to our knowledge, no study considered the two-dimensional Euclidean space model that was promoted by Suzuki and Yamashita [41]. A possible explanation for this situation is that the more complex the algorithm (or the system setting), the more difficult it becomes to rigorously find the worst possible execution.

3.1 Rendezvous Algorithms

We first quantify the maximum traveled distance and the average traveled distance for several known **Rendezvous** algorithms. We consider the *Center Of Gravity algorithm* [41], the two-color ASYNC algorithm (*Her2*) by Heriban et al. [30], the two-color algorithm (*Vig2*) by Viglietta [42], which is known to solve **Rendezvous** in SSYNC, and **Convergence** in ASYNC, the three-color algorithm (*Vig3*) by Viglietta [42], the four-color algorithm (*Das4*) by Das et al. [16, 17]. We also investigate the algorithms assuming unreliable compasses by Izumi et al. [33]: the SSYNC static-error compass algorithm (*Stat SSYNC*), which, despite its name, works in ASYNC, the SSYNC dynamic-error compass algorithm (*Dyn SSYNC*), which does not work in ASYNC, and the ASYNC static-error compass algorithm (*Dyn ASYNC*).

We take advantage of the modularity of our simulator. The `Robot` class now carries several new properties: `color`, the color a robot presently displays ; `compass`, the type of compass and error, *i.e.* 'none', 'static' or 'dynamic' ; `compass_error`, the maximum error allowed for the compass ; and `compass_offset`, the current compass error. The color is changed at the end of the `COMPUTE` method. Depending on the value of `compass`, `compass_offset` is either chosen during the initialization, or at the beginning of every `LOOK` method.

Each algorithm is first carefully analyzed on paper to find the worst possible execution. Simulations are then run according to the aforementioned protocols. Due to limitations described in Section 2.5, we actually assess those protocols for a degraded notion of **Convergence** rather than **Gathering**. The distance traveled is expressed relatively to the initial distance between the two robots. In practice, the first robot is always located at $\{0,0\}$ and the second robot is placed at random on the circle of radius 1 centered on $\{0,0\}$. Algorithms are tested with no initial pending moves, as arbitrary pending moves would render fuel efficiency mostly impossible to reliably monitor.

Results are summed up in Table 1. The red color denotes cases where the simulation was stuck in non-gathered cycles, and had to be manually unstuck. Details as to why this happened are provided below.

For scale, running 4 instances of *Vig3* for one hour under the ASYNC scheduler resulted in $\simeq 14$ million total individual executions.

	Rigid FSYNC	Rigid SSYNC	Rigid ASYNC	Non-Rigid FSYNC	Non-Rigid SSYNC	Non-Rigid ASYNC
CoG	1 / 1	1 / 1	2.66 / 3	1 / 1	1 / 1	2.36 / 3
Das4	1 / 1	1 / 1	1 / 1	1 / 1	1 / 1	1 / 1
Vig3	1 / 1	1 / 1	2 / 1	1 / 1	1 / 1	2 / 1
Vig2	1 / 1	1 / 1	2.94 / 3	1 / 1	1 / 1	2.75 / 3
Her2	1 / 1	1 / 1	2 / 2	1 / 1	1 / 1	2.54 / 3
Stat SSYNC	Stuck	Stuck	Stuck	Stuck	Stuck	Stuck
Dyn SSYNC	1	1	N/A	1.41	1.41	N/A
Dyn ASYNC	6.2	6.2	6.2	49.7	43.4	44.4
Stat SSYNC ϵ	3	3	3	3.95	4	4
Dyn SSYNC ϵ	1	1	N/A	1.41	1.41	N/A
Dyn ASYNC ϵ	6.2	6.2	6.2	49.4	50.9	41

(a) Maximum Traveled Distance
Found / Predicted

	Rigid FSYNC	Rigid SSYNC	Rigid ASYNC	Non-Rigid FSYNC	Non-Rigid SSYNC	Non-Rigid ASYNC
CoG	1	1	1.35	1	1	1.08
Das4	1	1	1	1	1	1
Vig3	1	1	1.06	1	1	1.01
Vig2	1	1	1.12	1	1	1.04
Her2	1	1	1.15	1	1	1.03
Stat SSYNC	Stuck	Stuck	Stuck	Stuck	Stuck	Stuck
Dyn SSYNC	1	1	N/A	1.03	1.03	N/A
Dyn ASYNC	2.18	2.00	2.07	1.99	1.88	2.00
Stat SSYNC ϵ	1.95	1.59	1.80	1.43	1.34	1.46
Dyn SSYNC ϵ	1	1	N/A	1.03	1.03	N/A
Dyn ASYNC ϵ	2.21	1.93	2.09	1.96	1.91	2.04

(b) Average Traveled Distance

Table 1 – Maximum and Average Traveled Distances

While most results match the predictions, our pen and paper analysis missed a worst case execution for ASYNC *Vig3*, which was found by the simulator (highlighted in bold in Figure 1). This highlights the difficulty of manually finding the maximum distance even with simple algorithms and settings. It should be noted that rigid motion yields worst results than non-rigid. This is normal because increasing the traveled distance relies on picking a target outside of the $[r_1, r_2]$ segment, and when this is the case, performing the full motion increases the traveled distance more than performing it partially. Thus, unless stated otherwise, all further simulations assume rigid motion.

The difference between SSYNC and ASYNC with respect to efficiency becomes apparent, as under the ASYNC scheduler, optimal fuel consumption mandates using four colors, while a simple oblivious algorithm is sufficient in SSYNC.

The algorithms using compasses yield the most interesting results. First, numerous simulations of the SSYNC static algorithm became stuck. These failures are due to the fact that the sine and cosine operations used in the algorithms tend to sum errors, and there is a possibility that a robot moves in a way that results in an angle of exactly 0, which actually randomly yields an angle of either $0 - \epsilon$ or $0 + \epsilon$, where ϵ is a very small positive number. This in turn results in unsolvable cycles that prevent **Convergence**. As ϵ was never larger than 10^{-9} , we chose to prevent this behavior by slightly enlarging the interval of the condition that should be triggered on an angle of zero to an angle in $[-10^{-6}, 10^{-6}]$. We do the same for all conditions for consistency. So any condition that should be true for angles in $[A, B[$ are now true for angles in $[A - 10^{-6}, B - 10^{-6}[$, in $[A, B]$ now in $[A - 10^{-6}, B + 10^{-6}]$, in $]A, B]$ now in $]A + 10^{-6}, B + 10^{-6}]$ and in $]A, B[$ now in $]A + 10^{-6}, B - 10^{-6}[$.

Interestingly, this new condition only had notable impact on the static error algorithm. Indeed, these errors could be seen as small dynamic random angle errors. Since the static error algorithm is not designed to be resilient against dynamic errors, it fails whenever they appear.

This also demonstrate the resilience of the dynamic error algorithms.

3.2 Convergence For n Robots

Cohen and Peleg [12] proved the Center of Gravity (CoG) algorithm solves **Convergence** for n robots under the ASYNC scheduler. We analyze the fuel consumption of the algorithm under both the SSYNC and ASYNC schedulers. Results for the minimum, maximum, and average distance traveled are show in table 2. We use the sum of the distances to the CoG in the initial configuration as a baseline unit of distance, *i.e.* the distance traveled in FSYNC.

	Minimum	Maximum	Average
ASYNC			
n = 2	1.00	2.74	1.34
n = 3	0.771	2.96	1.39
n = 4	0.751	2.82	1.37
n = 5	0.818	2.69	1.35
n = 10	1.00	1.97	1.26
n = 25	NC	NC	NC
n = 50	NC	NC	NC
SSYNC			
n = 2	1.00	1.00	1.00
n = 3	0.752	1.50	1.08
n = 4	0.689	1.91	1.10
n = 5	0.711	2.18	1.11
n = 10	0.883	2.15	1.12
n = 25	0.977	1.57	1.09
n = 50	0.995	1.31	1.07

Table 2 – Traveled Distances for CoG

It should be noted that, while previous results are based on at least hundreds of thousands of simulations, due to the increase in simulation complexity, in ASYNC, for $n = 25$, only 31 simulations could be computed under an hour. So they were discarded. Similarly, for $n = 50$, no simulation could be finished under an hour.

Looking at the results, one element immediately jumps out: for $n \geq 3$, the CoG algorithm wastes movements. This is easy to understand: robots move towards the center of gravity, which for 3 or more robots is different from the geometric median (*a.k.a.* the Weber point), which would actually minimize movement. Our tests seem to indicate that aiming for the median instead of the CoG can reduce traveled distance by up to 30%. However, it is a known result that no explicit formula for the geometric median exists. As a result, in practice, when trying to minimize traveled distance, **Convergence** for n robots should rely on an approximation of the geometric median rather than the center of gravity.

4 Analyzing Algorithms in Realistic Settings

In Section 3, the simulation of inaccurate compasses yielded extremely interesting results. To follow this track, we now focus in this section on the setting where sensors are inaccurate. In more details, we analyze the Center of Gravity (CoG) algorithm for **Rendezvous** in this setting, as well as the **Geoleader Election** algorithm by Canepa and Gradinariu Potop-Butucaru [10].

4.1 Visibility Sensor Errors

To study the impact of inaccurate sensors, we consider three different models for vision error. For a robot r_1 looking at a robot r_2 located in (x, y) in the Cartesian coordinate system centered

at r_1 , and located at (r, θ) in the polar coordinate system centered at r_1 , we define:

- The *absolute* error model [35] uses a constant value err . A first number R_{err} is picked uniformly at random in $[0, err]$, and a second θ_{err} in $[0, 2\pi]$. The perceived position of r_2 is then $(x + R_{err}\cos(\theta_{err}), y + R_{err}\sin(\theta_{err}))$.
- The *relative* error model [13] uses two constants err_{dist} and err_{angle} . Two numbers R_{err} and θ_{err} are picked uniformly at random in $[-err_{dist}, err_{dist}]$ and $[-err_{angle}, err_{angle}]$. The polar coordinates of r_2 are then perceived to be $(r + r * R_{err}, \theta + \theta_{err})$
- The *absolute-relative* error model is similar to relative error, but the perceived polar coordinates are $(r + R_{err}, \theta + \theta_{err})$

These error models are depicted in Figure 3.




Figure 3 – Types of Errors

The r_2 point is the actual location of robot r_2 . The red hashed area represents possible detected positions by robot r_1 .

It should be noted that each model could be used to accurately model errors for different types of sensors. The absolute error model is interesting because it is simple to compute, requires no change of coordinate system, uses a single parameter, and closely matches the behavior of robots where the LOOK phase is an abstraction of GPS-type coordinates exchanges [44]. The two relative models are more complex from a computing perspective, but closely match the use

of either computer vision or telemetry sensors. Both carry an angular error matched with either proportional or absolute distance error. Which type of distance error is more appropriate would depend on the exact type of sensor.

These new error models drive adding three properties to the Robot class:

- `LOOK_error_type`, a string that defines the type of error and can be either 'none', 'relative', 'absolute', or 'abs-rel'.
- `LOOK_distance_error`, a float that matches either err or err_{dist} , depending on the type of error.
- `LOOK_angle_error`, a float that matches err_{angle} .

Robots then chose the corresponding error (with parameters chosen uniformly at random) when performing their LOOK operation.

4.2 Convergence for $n=2$

Convergence with vision error using the CoG algorithm has already been studied by Cohen and Peleg [13]. The error model they considered is identical to our relative error model. Their paper states that **Convergence** with distance error using the CoG algorithm is impossible in the general case. This is, however, only true for $n \geq 3$, which the authors omit to mention. In the case $n = 2$, it appears to be theoretically impossible to make the algorithm diverge for a distance error smaller than a 100%, or $err = 1$. We can reasonably ignore the case of an error greater than 100%, as it would allow for a robot to perceive another one directly behind itself.

To our knowledge, no formal result exists regarding the angle error. In theory, the maximum angle error is π . We simulate **Convergence** for $n = 2$ robots using the CoG algorithm for the relative error model. The error for each robot is chosen uniformly at random at the beginning of the execution.

We must also consider the now possible case of a diverging algorithm. Since the execution is random, any setting should *eventually* converge. However, we must put a reasonable stopping condition in case the execution is clearly diverging. We chose to activate the defeat condition if the distance between the two robots becomes ten times larger than the distance in the initial configuration.

Note that the apparent decrease in maximum and average traveled distance for higher angle error is most likely due to the increase of diverging executions (fewer executions converge, but the traveled distance for those is shorter).

It appears clearly that the angular error has a much greater potential for both preventing **Convergence**, and making robots waste fuel. Indeed, when the angular error remains below $3\pi/5$, a distance error up to 100% can be tolerated with no performance loss.

To give some perspective, the realistic setting of a 10% vision error with a 1° angle error yields a maximum traveled distance of 1.221 and an average of 1.036, with no divergent executions out of more than 500 million data points.

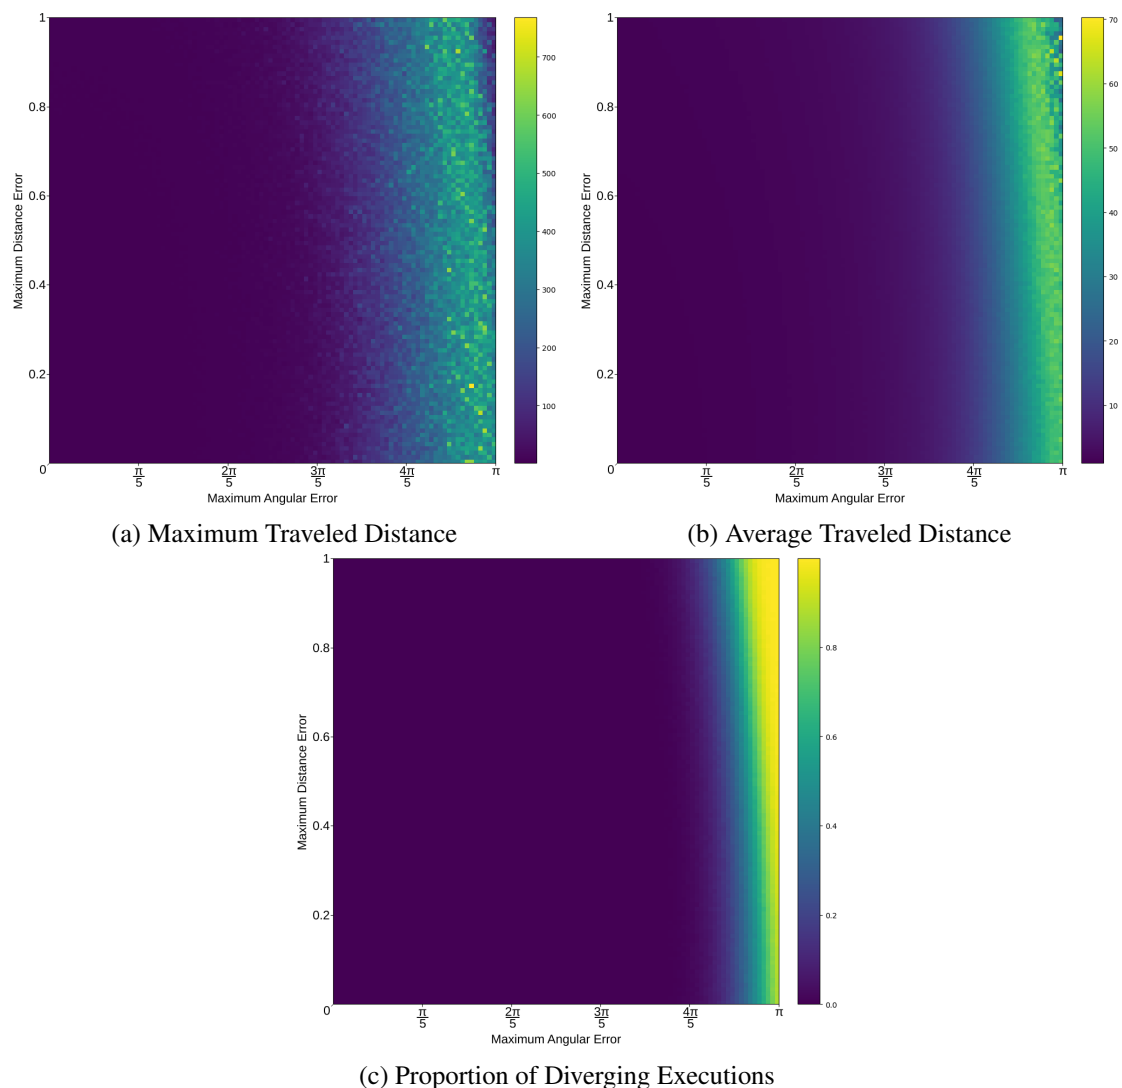


Figure 4 – Movement and Divergence of the CoG Algorithm for Two Robots with Inaccurate Visibility Sensors

4.3 Compass Errors

In the particular case of compass based algorithms of Izumi *et al.* [33], rendezvous is possible when the compasses are inaccurate. More specifically, the maximum tolerated errors are $\frac{\pi}{2}$, $\frac{\pi}{4}$ and $\frac{\pi}{6}$ for the static SSYNC, dynamic SSYNC, and dynamic ASYNC algorithms, respectively. In our simulation we chose static errors, for consistency, with values up to $\frac{49\pi}{100}$, $\frac{24\pi}{100}$ and $\frac{16\pi}{100}$, to avoid possible edge cases. Results of maximum and average traveled distances for these algorithms are detailed in Table 3.

We observe that the unreliable compasses are used in a way that makes robots rotate around each other until they are oriented in such a way that one robot moves while the other stays, regardless of the error. However, there are no provisions in these algorithms to limit distance increases during the rotating phases, which explains the results. Detailed observation shows the distance between the two robots can gradually diverge towards infinity during rotation and then converge to zero in a single cycle. This also demonstrated a problem for our **Convergence** criterion: robots could converge at rather large coordinates such that the coordinates of robots are in succession, but, since the accuracy of floating point numbers decreases as the number increase, the distance between the two robots was greater than 10^{-10} . As a result, we modified

	Rigid FSYNC	Rigid SSYNC	Rigid ASYNC	Non-Rigid FSYNC	Non-Rigid SSYNC	Non-Rigid ASYNC
Stat SSYNC	194	5.02E+03	4.73E+03	1.16E+06	5.33E+05	2.69E+07
Dyn SSYNC	5.88E+24	1.63E+13	N/A	4.69E+13	2.77E+11	N/A
Dyn ASYNC	6.21E+29	6.21E+29	6.21E+29	6.06E+127	5.54E+135	1.90E+123

(a) Maximum Traveled Distance

	Rigid FSYNC	Rigid SSYNC	Rigid ASYNC	Non-Rigid FSYNC	Non-Rigid SSYNC	Non-Rigid ASYNC
Stat SSYNC	6.14	6.34	10.7	8.35	7.99	54.3
Dyn SSYNC	1.61E+17	71.2	N/A	49.8	13.5	N/A
Dyn ASYNC	7.26E+23	7.75E+21	4.21E+25	1.54E+51	9.61E+38	8.60E+43


(b) Average Traveled Distance

Table 3 – Maximum and Average Traveled Distances for **Rendezvous** with Inaccurate Compasses

the criterion to $|r_1 r_2| < \max(10^{-10}, |Or_1| * 10^{-10})$, with O the point of coordinates $\{0, 0\}$.

4.4 Geoleader Election

Let us now consider **Geoleader Election** algorithm 2.3 by Canepa and Gradinariu Potop-Butucaru [10], for $n = 3$. Looking at our previous results from Section 2.5, we notice that the borders between each zone should be an issue for imperfect sensors, as different errors for different robots may lead to robots electing different **LEADER** robots.

Figure 5 – Example of **Leader Election** Failure Due to Imperfect Vision

We demonstrate how this phenomenon can occur in Figure 5 for the case of absolute vision error. On top is the actual configuration, where angles $\widehat{r_1 r_2 r_3}$ and $\widehat{r_2 r_1 r_3}$ are equal⁷, and angle $\widehat{r_1 r_3 r_2}$ is smaller than both, so r_3 should be elected. The red circle shows the possible perceived

⁷Because robots have no chirality, angles cannot reliably be distinguished from their opposite. So, two opposite angles may always be considered equal.

position of r_3 by r_1 and r_2 due to vision error. In the bottom left case, we show a possible perception by r_1 where r_1 should be elected **LEADER**, as $\widehat{r_2 r_1 r_3}$ is now greater than $\widehat{r_1 r_2 r_3}$. On the lower right, r_2 similarly thinks it should be elected. Now, two different robots consider themselves **LEADER** and the election process fails.

We now use the absolute model to simulate **Geoleader Election** with $err = 0.001$, for $n = 3$. This simulation yields $\simeq 0.1\%$ of errors in total, where two robots compute different **LEADER** robots, and is shown in figure 6.




Figure 6 – Simulation for 3-robot **Leader Election** with Absolute Vision Error
Yellow points represent configurations where the error generates two different **LEADER** robots.

5 Improved Convergence and Leader Election for Faulty Visibility Sensors

Following our observations of problematic behaviors in Sections 3 and 4, we provide two new algorithms: a fuel efficient **Convergence** algorithm for two robots, and a **Geoleader Election** algorithm that is resilient to faulty visibility sensors.

5.1 Fuel Efficient Convergence

We provide a new algorithm (Algorithm 5.1) for the **ASYN**C **Convergence** of two robots. Our algorithm is a simplified version of the two-color algorithm by Viglietta [42], which does *not* solve **Gathering** (while Viglietta’s algorithm does solve **Gathering** in **SSYN**C). Our algorithm however ensures that no target can ever be outside of the segment between the two robots, ensuring no wasted moves, and that there exists a scheduling such that convergence is eventually achieved. It is denoted by **FEC** (Fuel Efficient **Convergence**, presented in Figure 7). Our algorithm still uses two colors (**BLACK** and **WHITE**), and when observing the other robot’s color, the observing robot either remains still (the ‘Self’ target) or goes to the computed midpoint between the two robots (the ‘Midpoint’ target), possibly switching its color to the opposite one.

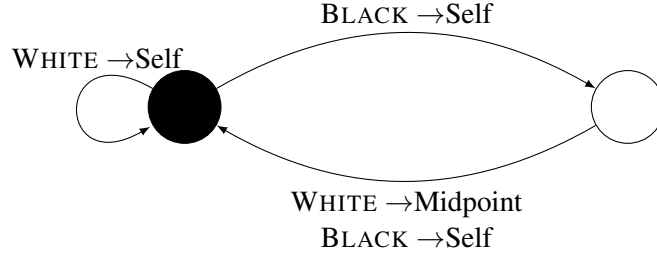


Figure 7 – FEC: Fuel Efficient **Convergence** Algorithm for Two Robots

Algorithm 5.1 FEC: Fuel Efficient **Convergence** Algorithm for Two Robots

```

if me.color = WHITE
    me.color  $\Leftarrow$  BLACK
    if other.color = WHITE
        me.destination  $\Leftarrow$  other.position/2
else if me.color = BLACK
    if other.color = BLACK
        me.color  $\Leftarrow$  WHITE
  
```

As a sanity check, we ran this algorithm through our simulator for one hour (\simeq 30 million data points) under a randomized ASYNC scheduler and could not find a single execution where the traveled distance was greater than the initial distance.

Theorem 5.1. *The Fuel Efficient **Convergence** Algorithm (5.1) guarantees the distance traveled for **Convergence** is never greater than the initial distance between the two robots under the ASYNC scheduler, assuming no pending moves in the initial configuration.*

Proof. First, we see that to achieve **Convergence** with an optimal distance, robots should always be moving towards each other. So, for robots to converge using more than the initial distance, it is required that, at one point in the execution, one robot moves *not towards* the other robot. We note that a network of two disoriented robots can be simplified as a line. In that sense, the only movement that can increase the maximum **Convergence** distance is when a robot moves opposite the other robot. In other words, when robots 'switch sides'. Let us now prove that no robot can target a robot while it is in its MOVE phase: Only the {WHITE,WHITE} snapshot can trigger a MOVE phase. Since this transition implies a change of color to BLACK at the end of the COMPUTE phase, robots that move can only be BLACK. So, if a robot is moving, it is BLACK and the other robot, regardless of color, cannot start moving because its snapshot is different from {WHITE,WHITE}. Furthermore, because robots switch to BLACK after moving, and can only switch to WHITE if the other robot is BLACK, no robot can execute multiple MOVE in sequence unless the other robot has executed at least a full cycle in between. So a robot cannot move multiple times while the other has pending moves. We look at what happens after each robot completes at least one full cycle. We assume r_1 performs a LOOK, and r_2 performs k cycles before r_1 finishes its MOVE. The distance after r_1 finishes its cycle is presented in table 4.

In the case of simultaneous M2H, the distance can be reduced down to zero, but robots cannot switch sides. In both other cases where a MOVE happens, the distance is reduced at most down to half, and robots cannot switch sides. Overall, in no cases can the robots move not towards one another, so the maximum distance traveled is always the initial distance between the two robots. \square


²As explained above, moving a second time requires at least a full cycle from the other robot.

	r_1 has a pending STAY	r_1 has a pending M2H
r_2 executes k STAY	X	$\left[\frac{X}{2}, X - \delta\right]$
r_2 executes 1 M2H ²	$\left[\frac{X}{2}, X - \delta\right]$	$[0, X - 2\delta]$

Table 4 – Distance after a full cycle of r_1 and k full cycles of r_2 with an initial distance of X

However, while the randomized scheduler we use for the simulator ensures convergence is always achieved, a rapid analysis of the algorithm shows that this algorithm ensures fuel efficiency, but does not actually ensure convergence. In fact, a simple SSYNC scheduling can infinitely prevent robots from moving. This further highlights that simulations and formal proofs are complementary. We conjecture that Fuel Efficient Convergence is not actually possible for two colors, and that algorithms using three colors may even yield Fuel Efficient Rendezvous (not just Convergence).

We also compare the resilience of this algorithm against vision errors with the center of gravity algorithm as a baseline in Figures 8a and 8b. Our results show that this algorithm is slightly more resilient to vision errors than CoG.



(a) Maximum Distance Traveled by CoG (top) and FEC (bottom)

(b) Average Distance Traveled by CoG (top) and FEC (bottom)

5.2 Error Resilient Geoleader Election

The **Geoleader Election** algorithm by Canepa and Gradinariu Potop-Butucaru [10] was *not* designed under the assumption that the visibility sensors could be prone to errors. In this subsection, we use this awareness to create a new, error-resilient, version of this algorithm, using our simulation framework.

Geoleader Election for Four Robots One intuitive way of building a fully resilient algorithm for **Leader Election** could be based on robots computing the bounds of the error zone. While this seems feasible for a 3-robot election, it becomes far less trivial for four robots or more. We present the results of a leader election for four robots in the appendix.

Proposed Algorithm In section 4, we used the simulation framework to detect failed elections caused by visibility sensor errors. Since mobile robots are able to run any algorithm during their COMPUTE phase, then they can also run the simulation framework to do precisely that.

The improved algorithm relies on the knowledge of the vision error model and its upper bounds to simulate random errors in a robot's position and snapshot and determine whether there exists a possibility of the other robots electing different **LEADER** robots.

Note that absolutely knowing that the election cannot fail (*i.e.*, the election cannot yield two different **LEADER** robots for two different robots) would require checking the entire surface of possible errors, which is not feasible in practice. So, we assume that robots perform a finite number of trials and decide accordingly.

Each robot internally simulates a position error for each robot in its snapshot within the known margins, performs a simulated election for each robot in its snapshot, and checks for discrepancies in the resulting **LEADER** robots. This is repeated with new random errors for a given number of tries, similar to a Monte-Carlo approach.

Once a robot believes the election process can succeed, it chooses the **LEADER** as in the original algorithm. Otherwise, it picks a random direction and distance, and performs a MOVE to "scramble" the network. This process repeats until all robots believe the election can succeed. This process is detailed in Algorithm 5.2.

Algorithm 5.2 Reliable Leader Election algorithm

```

L = self.COMPUTE('LeaderElection')
my_network = self.snapshot ∪ self
counter = 0
while counter < nb_tries do
  for r1 in my_network do
    rv = r1
    Change rv.x and rv.y randomly according to error parameters
    rv.snapshot = my_network / {r1}
    for r2 in rv.snapshot do
      Change r2.x and r2.y randomly according to error parameters
      Lv = rv.COMPUTE('LeaderElection')
      if L ≠ Lv
        Move randomly
        Exit
    counter += 1
  L is elected LEADER

```

We now perform simulations using this algorithm. Each point is sorted according to the following:

- If no robot detects a possible error, it is a valid point.
- If at least one robot has detected a possible error, and decided to move as a result, it is a detected possible error point.
- If no robot moves, but two robots have different **LEADER** robots, it is an undetected error point.

We measure the proportion of undetected error and possible error points for nb_{tries} between 0 and 30. Results are presented in Figure 9.




Figure 9 – Performance of the Error-Resilient Election Algorithm
 $err = 0.001$

Note that the number of undetected error points, while decreasing, does not reach zero under our testing conditions. Also, using a single internal simulation typically results in a $\sim 80\%$ reduction in the number of undetected error points. Using 10 internal simulations resulted in a reduction of 99.5% of undetected error points. Which number of internal simulations is the best suited would depend on both the speed of the leader election process and reliability of the obtained solution requirements.

Importantly, we notice that, were we to choose an error model and error bounds such that it models the possible errors of representing real numbers using limited precision floats, then this particular algorithm, when used with an infinitely large, similar to \mathbb{R}^2 number of random tries, can be made to reliably detect anomalies due to the errors induced by evolving in the continuous plane, yet only perceiving a discretized plane. Actually, we make the conjecture that this algorithm can be adapted to allow any algorithm that makes decisions based on robot locations to operate properly in a perceived discretized plane.

Furthermore, using this algorithm allows us to reduce the size of a snapshot to a finite, storable amount to realistically use the **SyncSim** protocol [17], and fully simulate the FSYNC scheduler in *LUMINOUS ASYNC*.

6 Conclusion

In this paper, we introduce a modular framework designed to simulate mobile robots for any given setting.

We discuss the limitations and constraints of this approach, and use it to compute the maximum distance traveled, or fuel efficiency, of multiple algorithms in several settings, with interesting results. In particular, we note that the algorithm by Izumi *et al.* [33] can lead to an unbounded increase in distance before eventually gathering. Similarly, the center of gravity algorithm is inherently sub-optimal for $n > 2$ robots, and robots should use an algorithm based on the geometric median instead.

We then use this framework to simulate inaccurate sensors for mobile robots and verify the behavior of **Convergence** and motion based **Leader Election** under this new model. We also introduce errors in the perception of colors for *LUMINOUS* robots performing state-of-the-art two-robot **Gathering**.

Finally, we designed two new algorithms. The first one is designed to perform two-robot **Convergence** under the ASYNC scheduler with optimal fuel efficiency. The second algorithm uses the simulator itself to allow robots to solve motion based **Leader Election** with inaccurate sensors. The latter can be adapted to allow for decision making algorithm, such as **Leader Election**, to function using discretized snapshots, and so, to use the **SyncSim** protocol to simulate the FSYNC scheduler in *LUMINOUS ASYNC*.

Overall, this framework achieves its planned objective of being both easy to use and able to produce useful results for researchers. As a test, we timed the full implementation, and testing in FSYNC, SSYNC and ASYNC, of the two color **Rendezvous** algorithm from Viglietta to require less than half an hour, including basic network monitoring and testing.

The source code and instructions for our simulator are provided in the appendix and at the following repository: <https://github.com/UberPanda/PyBlot-Sim>

Future Work

As we already discussed, our simulator is modular to allow for use for any given algorithm and model. So it seems logical that it should, ideally, implement every existing model and test all major algorithms in the literature, such as mutual visibility for opaque robots.

Furthermore, while interesting for researchers, our simulator is not a tool for formal proofs. However, one could also argue that in its current state, we have not proven that the simulator actually simulates mobile robots, even within our degraded hypotheses. We believe that the simulator itself should be formally proven to match the model of mobile robots it claims to simulate. Note that the usefulness of this proof would be limited, as the addition of any new module may require proving the entire simulator again.

Finally, our **Leader Election** algorithm for errors in vision is able to function in a continuous setting using discretized snapshots. The design philosophy behind this algorithm of using randomized tries to simulate sensor errors is not specific to the **Leader Election** problem, and it could be used for other algorithms that rely on making decisions based on the locations of robots in the network and that are sensitive to errors in perception. Building new algorithms that can use these finite snapshots allows us to use the **SyncSim** protocol [17] and simulate a FSYNC scheduler in *LUMINOUS ASYNC* and would be a major advantage for resilience to asynchrony.

References

- [1] Hideki Ando, Yoshinobu Oasa, Ichiro Suzuki, and Masafumi Yamashita. “Distributed memoryless point convergence algorithm for mobile robots with limited visibility”. In: *IEEE Trans. Robotics and Automation* 15.5 (1999), pp. 818–828. DOI: 10.1109/70.795787.
- [2] Divansh Arora, Parikshit Maini, Pedro Pinacho Davidson, and Christian Blum. “Route planning for cooperative air-ground robots with fuel constraints: an approach based on CMSA”. In: *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2019, Prague, Czech Republic, July 13-17, 2019*, ed. by Anne Auger et al. ACM, 2019, pp. 207–214. DOI: 10.1145/3321707.3321820.
- [3] Cédric Auger, Zohir Bouzid, Pierre Courtieu, Sébastien Tixeuil, and Xavier Urbain. “Certified Impossibility Results for Byzantine-Tolerant Mobile Robots”. In: *Stabilization, Safety, and Security of Distributed Systems - 15th International Symposium, SSS 2013, Osaka, Japan, November 13-16, 2013. Proceedings*, ed. by Teruo Higashino et al. Vol. 8255. Lecture Notes in Computer Science. Springer, 2013, pp. 178–190. DOI: 10.1007/978-3-319-03089-0_13.
- [4] Thibaut Balabonski, Pierre Courtieu, Robin Pelle, Lionel Rieg, Sébastien Tixeuil, and Xavier Urbain. “Continuous vs. Discrete Asynchronous Moves: A Certified Approach for Mobile Robots”. In: *Networked Systems - 7th International Conference, NETYS 2019, Marrakech, Morocco, June 19-21, 2019, Revised Selected Papers*, ed. by Mohamed Faouzi Atig et al. Vol. 11704. Lecture Notes in Computer Science. Springer, 2019, pp. 93–109. DOI: 10.1007/978-3-030-31277-0_7.
- [5] Thibaut Balabonski, Amélie Delga, Lionel Rieg, Sébastien Tixeuil, and Xavier Urbain. “Synchronous Gathering without Multiplicity Detection: a Certified Algorithm”. In: *Theory Comput. Syst.* 63.2 (2019), pp. 200–218. DOI: 10.1007/s00224-017-9828-z.
- [6] Thibaut Balabonski, Robin Pelle, Lionel Rieg, and Sébastien Tixeuil. “A Foundational Framework for Certified Impossibility Results with Mobile Robots on Graphs”. In: *Proceedings of the 19th International Conference on Distributed Computing and Networking, ICDCN 2018, Varanasi, India, January 4-7, 2018*, ed. by Paolo Bellavista et al. ACM, 2018, 5:1–5:10. DOI: 10.1145/3154273.3154321.
- [7] Béatrice Bérard, Pierre Courtieu, Laure Millet, Maria Potop-Butucaru, Lionel Rieg, Nathalie Sznajder, Sébastien Tixeuil, and Xavier Urbain. “[Invited Paper] Formal Methods for Mobile Robots: Current Results and Open Problems”. In: *International Journal of Informatics Society* 7.3 (2015), pp. 101–114.
- [8] Béatrice Bérard, Pascal Lafourcade, Laure Millet, Maria Potop-Butucaru, Yann Thierry-Mieg, and Sébastien Tixeuil. “Formal verification of mobile robot protocols”. In: *Distributed Comput.* 29.6 (2016), pp. 459–487. DOI: 10.1007/s00446-016-0271-1.
- [9] François Bonnet, Xavier Défago, Franck Petit, Maria Potop-Butucaru, and Sébastien Tixeuil. “Discovering and Assessing Fine-Grained Metrics in Robot Networks Protocols”. In: *33rd IEEE International Symposium on Reliable Distributed Systems Workshops, SRDS Workshops 2014, Nara, Japan, October 6-9, 2014*. IEEE Computer Society, 2014, pp. 50–59. DOI: 10.1109/SRDSW.2014.34.
- [10] Davide Canepa and Maria Gradinariu Potop-Butucaru. “Stabilizing Flocking Via Leader Election in Robot Networks”. In: *Stabilization, Safety, and Security of Distributed Systems, 9th International Symposium, SSS 2007, Paris, France, November 14-16, 2007, Proceedings*, ed. by Toshimitsu Masuzawa et al. Vol. 4838. Lecture Notes in Computer Science. Springer, 2007, pp. 52–66. ISBN: 978-3-540-76626-1. DOI: 10.1007/978-3-540-76627-8_7.

- [11] Jérémie Chalopin, Shantanu Das, Matús Mihalák, Paolo Penna, and Peter Widmayer. “Data Delivery by Energy-Constrained Mobile Agents”. In: *Algorithms for Sensor Systems - 9th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics, ALGOSENSORS 2013, Sophia Antipolis, France, September 5-6, 2013, Revised Selected Papers*, ed. by Paola Flocchini et al. Vol. 8243. Lecture Notes in Computer Science. Springer, 2013, pp. 111–122. ISBN: 978-3-642-45345-8. DOI: 10.1007/978-3-642-45346-5_9.
- [12] Reuven Cohen and David Peleg. “Convergence Properties of the Gravitational Algorithm in Asynchronous Robot Systems”. In: *SIAM J. Comput.* 34.6 (2005), pp. 1516–1528. DOI: 10.1137/S0097539704446475.
- [13] Reuven Cohen and David Peleg. “Convergence of Autonomous Mobile Robots with Inaccurate Sensors and Movements”. In: *SIAM J. Comput.* 38.1 (2008), pp. 276–302. DOI: 10.1137/060665257.
- [14] Pierre Courtieu, Lionel Rieg, Sébastien Tixeuil, and Xavier Urbain. “Impossibility of gathering, a certification”. In: *Inf. Process. Lett.* 115.3 (2015), pp. 447–452. DOI: 10.1016/j.ip1.2014.11.001.
- [15] Pierre Courtieu, Lionel Rieg, Sébastien Tixeuil, and Xavier Urbain. “Certified Universal Gathering in \mathbb{R}^2 for Oblivious Mobile Robots”. In: *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016. Proceedings*, ed. by Cyril Gavoille et al. Vol. 9888. Lecture Notes in Computer Science. Springer, 2016, pp. 187–200. DOI: 10.1007/978-3-662-53426-7_14.
- [16] Shantanu Das, Paola Flocchini, Giuseppe Prencipe, Nicola Santoro, and Masafumi Yamashita. “The Power of Lights: Synchronizing Asynchronous Robots Using Visible Bits”. In: *2012 IEEE 32nd International Conference on Distributed Computing Systems, Macau, China, June 18-21, 2012*. IEEE Computer Society, 2012, pp. 506–515. ISBN: 978-1-4577-0295-2. DOI: 10.1109/ICDCS.2012.71.
- [17] Shantanu Das, Paola Flocchini, Giuseppe Prencipe, Nicola Santoro, and Masafumi Yamashita. “Autonomous mobile robots with lights”. In: *Theor. Comput. Sci.* 609 (2016), pp. 171–184. DOI: 10.1016/j.tcs.2015.09.018.
- [18] Xavier Défago, Adam Heriban, Sébastien Tixeuil, and Koichi Wada. “Brief Announcement: Model Checking Rendezvous Algorithms for Robots with Lights in Euclidean Space”. In: *33rd International Symposium on Distributed Computing, DISC 2019, October 14-18, 2019, Budapest, Hungary*, ed. by Jukka Suomela. Vol. 146. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, 41:1–41:3. DOI: 10.4230/LIPIcs.DISC.2019.41.
- [19] Xavier Défago, Adam Heriban, Sébastien Tixeuil, and Koichi Wada. “Using Model Checking to Formally Verify Rendezvous Algorithms for Robots with Lights in Euclidean Space”. In: *International Symposium on Reliable Distributed Systems, SRDS 2020, Shanghai, China, September 21-24, 2020*. IEEE, 2020, pp. 113–122. DOI: 10.1109/SRDS51746.2020.00019.
- [20] Stéphane Devismes, Anissa Lamani, Franck Petit, Pascal Raymond, and Sébastien Tixeuil. “Optimal Grid Exploration by Asynchronous Oblivious Robots”. In: *Stabilization, Safety, and Security of Distributed Systems - 14th International Symposium, SSS 2012, Toronto, Canada, October 1-4, 2012. Proceedings*, ed. by Andréa W. Richa et al. Vol. 7596. Lecture Notes in Computer Science. Springer, 2012, pp. 64–76. DOI: 10.1007/978-3-642-33536-5_7.
- [21] Yoann Dieudonné and Franck Petit. “Self-stabilizing gathering with strong multiplicity detection”. In: *Theor. Comput. Sci.* 428 (2012), pp. 47–57. DOI: 10.1016/j.tcs.2011.12.010.

- [22] Ha Thi Thu Doan, François Bonnet, and Kazuhiro Ogata. “Model Checking of a Mobile Robots Perpetual Exploration Algorithm”. In: *Structured Object-Oriented Formal Language and Method - 6th International Workshop, SOFL+MSVL 2016, Tokyo, Japan, November 15, 2016, Revised Selected Papers*, ed. by Shaoying Liu et al. Vol. 10189. Lecture Notes in Computer Science. 2016, pp. 201–219. DOI: 10.1007/978-3-319-57708-1_12.
- [23] Ha Thi Thu Doan, François Bonnet, and Kazuhiro Ogata. “Model Checking of Robot Gathering”. In: *21st International Conference on Principles of Distributed Systems, OPODIS 2017, Lisbon, Portugal, December 18-20, 2017*, ed. by James Aspnes et al. Vol. 95. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 12:1–12:16. DOI: 10.4230/LIPIcs.OPODIS.2017.12.
- [24] Mirosław Dynia, Mirosław Korzeniowski, and Christian Schindelhauer. “Power-Aware Collective Tree Exploration”. In: *Architecture of Computing Systems - ARCS 2006, 19th International Conference, Frankfurt/Main, Germany, March 13-16, 2006, Proceedings*, ed. by Werner Grass et al. Vol. 3894. Lecture Notes in Computer Science. Springer, 2006, pp. 341–351. ISBN: 3-540-32765-7. DOI: 10.1007/11682127_24.
- [25] Paola Flocchini. “Gathering”. In: *Distributed Computing by Mobile Entities, Current Research in Moving and Computing*, ed. by Paola Flocchini et al. Vol. 11340. Lecture Notes in Computer Science. Springer, 2019, pp. 63–82. DOI: 10.1007/978-3-030-11072-7_4.
- [26] Paola Flocchini, Giuseppe Prencipe, and Nicola Santoro. *Distributed Computing by Oblivious Mobile Robots*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2012. DOI: 10.2200/S00440ED1V01Y201208DCT010.
- [27] Paola Flocchini, Giuseppe Prencipe, Nicola Santoro, and Peter Widmayer. “Gathering of asynchronous robots with limited visibility”. In: *Theor. Comput. Sci.* 337.1-3 (2005), pp. 147–168. DOI: 10.1016/j.tcs.2005.01.001.
- [28] Noam Gordon, Yotam Elor, and Alfred M. Bruckstein. “Gathering Multiple Robotic Agents with Crude Distance Sensing Capabilities”. In: *Ant Colony Optimization and Swarm Intelligence, 6th International Conference, ANTS 2008, Brussels, Belgium, September 22-24, 2008. Proceedings*, ed. by Marco Dorigo et al. Vol. 5217. Lecture Notes in Computer Science. Springer, 2008, pp. 72–83. ISBN: 978-3-540-87526-0. DOI: 10.1007/978-3-540-87527-7_7.
- [29] Noam Gordon, Israel A. Wagner, and Alfred M. Bruckstein. “Gathering Multiple Robotic A(ge)nts with Limited Sensing Capabilities”. In: *Ant Colony Optimization and Swarm Intelligence, 4th International Workshop, ANTS 2004, Brussels, Belgium, September 5 - 8, 2004, Proceedings*, ed. by Marco Dorigo et al. Vol. 3172. Lecture Notes in Computer Science. Springer, 2004, pp. 142–153. ISBN: 3-540-22672-9. DOI: 10.1007/978-3-540-28646-2_13.
- [30] Adam Heriban, Xavier Défago, and Sébastien Tixeuil. “Optimally Gathering Two Robots”. In: *Proceedings of the 19th International Conference on Distributed Computing and Networking, ICDCN 2018, Varanasi, India, January 4-7, 2018*, ed. by Paolo Bellavista et al. ACM, 2018, 3:1–3:10. DOI: 10.1145/3154273.3154323.
- [31] Adam Heriban and Sébastien Tixeuil. “Mobile Robots with Uncertain Visibility Sensors”. In: *Structural Information and Communication Complexity - 26th International Colloquium, SIROCCO 2019, L'Aquila, Italy, July 1-4, 2019, Proceedings*, ed. by Keren Censor-Hillel et al. Vol. 11639. Lecture Notes in Computer Science. Springer, 2019, pp. 349–352. ISBN: 978-3-030-24921-2. DOI: 10.1007/978-3-030-24922-9_27.

- [32] Anthony Honorat, Maria Potop-Butucaru, and Sébastien Tixeuil. “Gathering fat mobile robots with slim omnidirectional cameras”. In: *Theor. Comput. Sci.* 557 (2014), pp. 1–27. DOI: 10.1016/j.tcs.2014.08.004.
- [33] Taisuke Izumi, Samia Souissi, Yoshiaki Katayama, Nobuhiro Inuzuka, Xavier Défago, Koichi Wada, and Masafumi Yamashita. “The Gathering Problem for Two Oblivious Robots with Unreliable Compasses”. In: *SIAM J. Comput.* 41.1 (2012), pp. 26–46. DOI: 10.1137/100797916.
- [34] Giuseppe Antonio Di Luna, Paola Flocchini, Federico Poloni, Nicola Santoro, and Giovanni Viglietta. “The Mutual Visibility Problem for Oblivious Robots”. In: *Proceedings of the 26th Canadian Conference on Computational Geometry, CCCG 2014, Halifax, Nova Scotia, Canada, 2014*. Carleton University, Ottawa, Canada, 2014. URL: <http://www.cccg.ca/proceedings/2014/papers/paper51.pdf>.
- [35] Sonia Martínez. “Practical multiagent rendezvous through modified circumcenter algorithms”. In: *Automatica* 45.9 (2009), pp. 2010–2017. DOI: 10.1016/j.automatica.2009.05.013.
- [36] Nicholas Metropolis and S. Ulam. “The Monte Carlo Method”. In: *Journal of the American Statistical Association* 44.247 (1949), pp. 335–341.
- [37] Laure Millet, Maria Potop-Butucaru, Nathalie Sznajder, and Sébastien Tixeuil. “On the Synthesis of Mobile Robots Algorithms: The Case of Ring Gathering”. In: *Stabilization, Safety, and Security of Distributed Systems - 16th International Symposium, SSS 2014, Paderborn, Germany, September 28 - October 1, 2014. Proceedings*, ed. by Pascal Felber et al. Vol. 8756. Lecture Notes in Computer Science. Springer, 2014, pp. 237–251. DOI: 10.1007/978-3-319-11764-5_17.
- [38] Sasha Rubin, Florian Zuleger, Aniello Murano, and Benjamin Aminof. “Verification of Asynchronous Mobile-Robots in Partially-Known Environments”. In: *PRIMA 2015: Principles and Practice of Multi-Agent Systems - 18th International Conference, Bertinoro, Italy, October 26-30, 2015, Proceedings*, ed. by Qingliang Chen et al. Vol. 9387. Lecture Notes in Computer Science. Springer, 2015, pp. 185–200. DOI: 10.1007/978-3-319-25524-8_12.
- [39] Arnaud Sangnier, Nathalie Sznajder, Maria Potop-Butucaru, and Sébastien Tixeuil. “Parameterized verification of algorithms for oblivious robots on a ring”. In: *2017 Formal Methods in Computer Aided Design, FMCAD 2017, Vienna, Austria, October 2-6, 2017*, ed. by Daryl Stewart et al. IEEE, 2017, pp. 212–219. DOI: 10.23919/FMCAD.2017.8102262.
- [40] Arnaud Sangnier, Nathalie Sznajder, Maria Potop-Butucaru, and Sébastien Tixeuil. “Parameterized verification of algorithms for oblivious robots on a ring”. In: *Formal Methods Syst. Des.* 56.1 (2020), pp. 55–89. DOI: 10.1007/s10703-019-00335-y.
- [41] Ichiro Suzuki and Masafumi Yamashita. “Distributed Anonymous Mobile Robots: Formation of Geometric Patterns”. In: *SIAM J. Comput.* 28.4 (1999), pp. 1347–1363. DOI: 10.1137/S009753979628292X.
- [42] Giovanni Viglietta. “Rendezvous of Two Robots with Visible Bits”. In: *Algorithms for Sensor Systems - 9th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics, ALGOSENSORS 2013, Sophia Antipolis, France, September 5-6, 2013, Revised Selected Papers*, ed. by Paola Flocchini et al. Vol. 8243. Lecture Notes in Computer Science. Springer, 2013, pp. 291–306. DOI: 10.1007/978-3-642-45346-5_21.

- [43] Kenta Yamamoto, Taisuke Izumi, Yoshiaki Katayama, Nobuhiro Inuzuka, and Koichi Wada. “The optimal tolerance of uniform observation error for mobile robot convergence”. In: *Theor. Comput. Sci.* 444 (2012), pp. 77–86. DOI: 10.1016/j.tcs.2012.04.038.
- [44] Rami Yared, Xavier Défago, Julien Iguchi-Cartigny, and Matthias Wiesmann. “Collision Prevention Platform for a Dynamic Group of Asynchronous Cooperative Mobile Robots”. In: *J. Networks* 2.4 (2007), pp. 28–39. DOI: 10.4304/jnw.2.4.28-39.
- [45] Chanyeol Yoo, Robert Fitch, and Salah Sukkarieh. “Online task planning and control for fuel-constrained aerial robots in wind fields”. In: *Int. J. Robotics Res.* 35.5 (2016), pp. 438–453. DOI: 10.1177/0278364915595278.

A Appendix: example of an Instance of the Simulator

We present a minimum working example of an instance of the simulator. It simulates executions of the Vig2 algorithm [42] in the standard *OBLLOT* model with rigid motion, under the FSYNC scheduler. It monitors the number of cycles needed to complete degraded gathering⁸. Results predictably show the possible need of two full cycles in the case where both robots start in the BLACK color. For better readability, more advanced features of the simulator are not included.

Listing 1 – Robot Class File: Common/lib_robot.py

```
import Common.lib_algorithms as lib_algos

class Robot:
    def __init__(self, name, x, y, color=0):
        self.name = name
        self.x = x ## Real position in the network
        self.y = y

        self.phase = 'WAITING'

        # Available info for COMPUTE :

        self.snapshot = [] ## snapshot
        self.color = color ## color
        self.target = ()

    def LOOK(self, network, scheduler):
        self.snapshot = []
        for R2 in network:
            if self != R2:
                R2x = R2.x
                R2y = R2.y
                self.snapshot.append(Robot(R2.name, R2x, R2y, R2.color))

        self.phase = 'COMPUTING'

    def COMPUTE(self, algo):
        try:
            result = getattr(lib_algos, algo)(self)
        except:
            raise AttributeError('This algorithm does not exist in the current version (or you
                ↪ made a typo ?)')

        return result

    def MOVE(self):
        self.x = self.target[0]
        self.y = self.target[1]

        self.target = ()
        self.phase = 'WAITING'
```

Listing 2 – Miscellaneous Functions File: Common/lib_misc_functions.py

```
from math import sqrt
```

⁸i.e. robots are closer than 10^{-10} with the initial distance being 1.

```

def rob_dist(R1,R2): # Using Robots
    return dist((R1.x,R1.y),(R2.x,R2.y))

def dist(T1,T2): # Using Tuples
    return sqrt((T1[0] - T2[0])**2+(T1[1] - T2[1])**2)

```

Listing 3 – Algorithms File: Common/lib_algorithms.py

```

def vig2(R1):
    R2 = R1.snapshot[0]
    if R1.color == 0:
        if R2.color == 0:
            R1.target = ((R1.x+R2.x)/2, (R1.y+R2.y)/2)
            R1.color = 1
            R1.phase = 'MOVING'
        else:
            R1.target = (R2.x, R2.y)
            R1.phase = 'MOVING'
    else:
        R1.phase = 'WAITING'
        if R2.color == 1:
            R1.color = 0

```

Listing 4 – Scheduler File: Common/lib_schedulers.py

```

def scheduler(sched,network,algo):

    if sched == 'FSYNC':
        for R1 in network:
            R1.LOOK(network,'FSYNC')
        for R1 in network:
            R1.COMPUTE(algo)
        for R1 in network:
            if R1.phase == 'MOVING':
                R1.MOVE()

    else:
        raise Exception('Unknown Scheduler')

```

Listing 5 – Simulation Functions File: Common/lib_sim_functions.py

None Needed for this Example

Listing 6 – Simulation File: Vig2.py

```

from random import SystemRandom, choice, uniform
from math import sqrt, pi, cos, sin
from Common.lib_robot import Robot
from Common.lib_schedulers import scheduler
from Common.lib_misc_functions import rob_dist

_sysrand = SystemRandom()

#####

SSTAB = True
COLOR_RANGE = range(2) # Total number of colors.

#####

```



```

simus = 0
act_list = []

while simus < 1000000:
    simus += 1

    # Network initialization

    network = []
    network.append(Robot('r1',0,0,0))

    ang = uniform(-pi,pi)
    network.append(Robot('r2',cos(ang),sin(ang),0))

    for R1 in network:
        if SSTAB == True:
            R1.color = choice(COLOR_RANGE)

    ## Beginning of the simulation

    steps = 0
    while True:

        scheduler('FSYNC',network,'vig2')

        ## Live state monitoring

        steps += 1

        ## Victory condition

        if rob_dist(network[0],network[1]) < 10**(-10):
            break

        ## Defeat condition

    act_list.append(steps)

print('FSYNC R Vig2')
print('Min : ' + str(min(act_list)))
print('Max : ' + str(max(act_list)))
print('Avg : ' + str(sum(act_list) / len(act_list)))

```

B Appendix: leader Election for 4 Robots

Figures 10 through 15 show the result of attempting to elect a **GEOLEADER** using algorithm B.1 by Canepa and Gradinariu Potop-Butucaru [10], which should fail whenever two robots are identically close to the center of the smallest enclosing circle. However, as we have shown, such cases are statistically impossible with perfect sensors and simply become a small subset of the error points of error-prone sensors.

Algorithm B.1 Original **Leader Election** Algorithm by Canepa and Gradinariu Potop-Butucaru [10] for Four or More Robots

```
Compute the smallest enclosing circle SEC
Compute the distance  $d_k$  to the center of SEC, for all robots  $1 \leq k \leq n$ 
if  $d_{myself} < d_k \forall k \neq myself$ , where  $1 \leq k \leq n$ 
    Become LEADER
    Exit
if  $d_{myself} \leq d_k \forall k \neq myself$ , where  $1 \leq k \leq n$ 
    Perform a Bernoulli trial with a probability of winning of  $p = \frac{1}{n}$ 
    if Trial won
        move a distance of  $d_{myself} \cdot p$  towards the center of the SEC
```

Robots r_1 and r_2 are fixed at coordinate $(-0.5, 0)$ and $(0.5, 0)$, respectively. Robot r_3 has a fixed location for each image, on a grid in the lower left quarter of the image. Symmetries of the network allow us to extrapolate results for the remainder of the positions of r_3 . The position of r_4 is chosen at uniformly at random, and each point show the result for a given position. The error is absolute with $err = 0.001$.

As before, colors red, green, and blue denote that robot r_1 , r_2 and r_3 are the chosen **GEOLEADER**, respectively. Color cyan denotes that robot r_4 is the chosen **GEOLEADER**, and yellow denotes that two different **GEOLEADER** robots have been elected due to sensor error. Similarly to previous simulations, each image contains one million points.

As can be seen from the following figures, computing precisely the bounds of the error zone is extremely complex in practice, as the formula would be different and more complex as the network size grows.

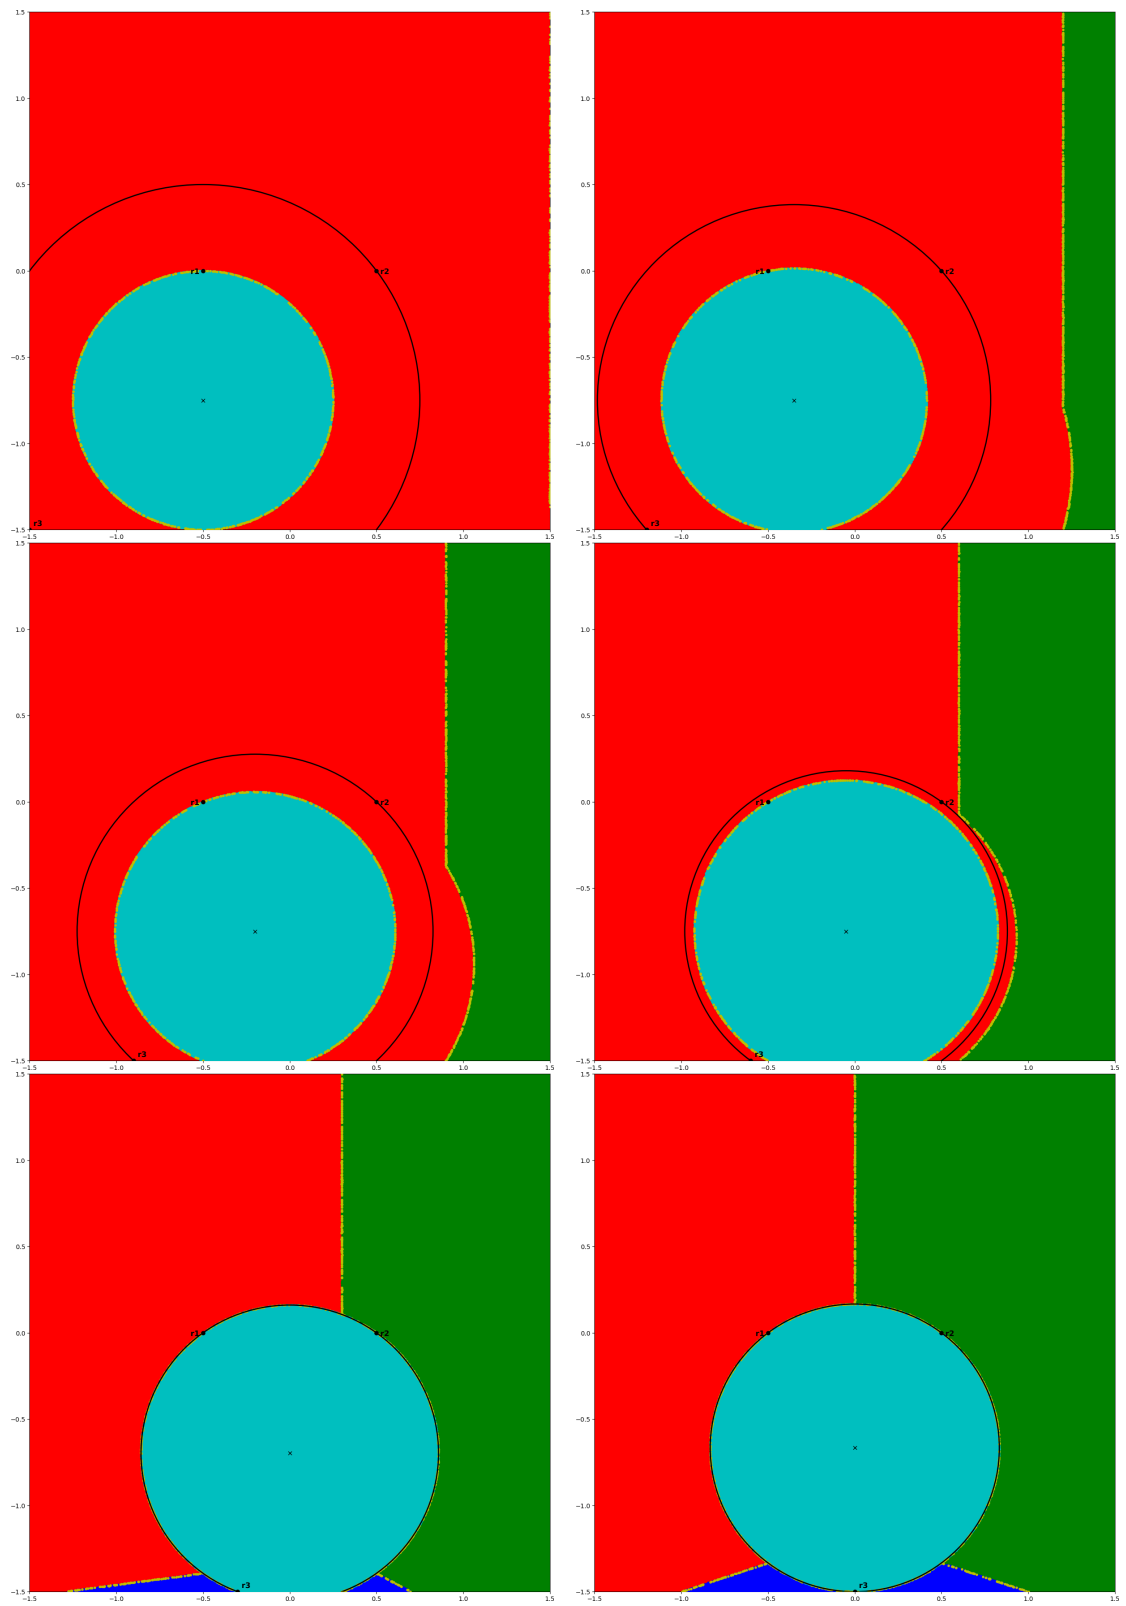


Figure 10

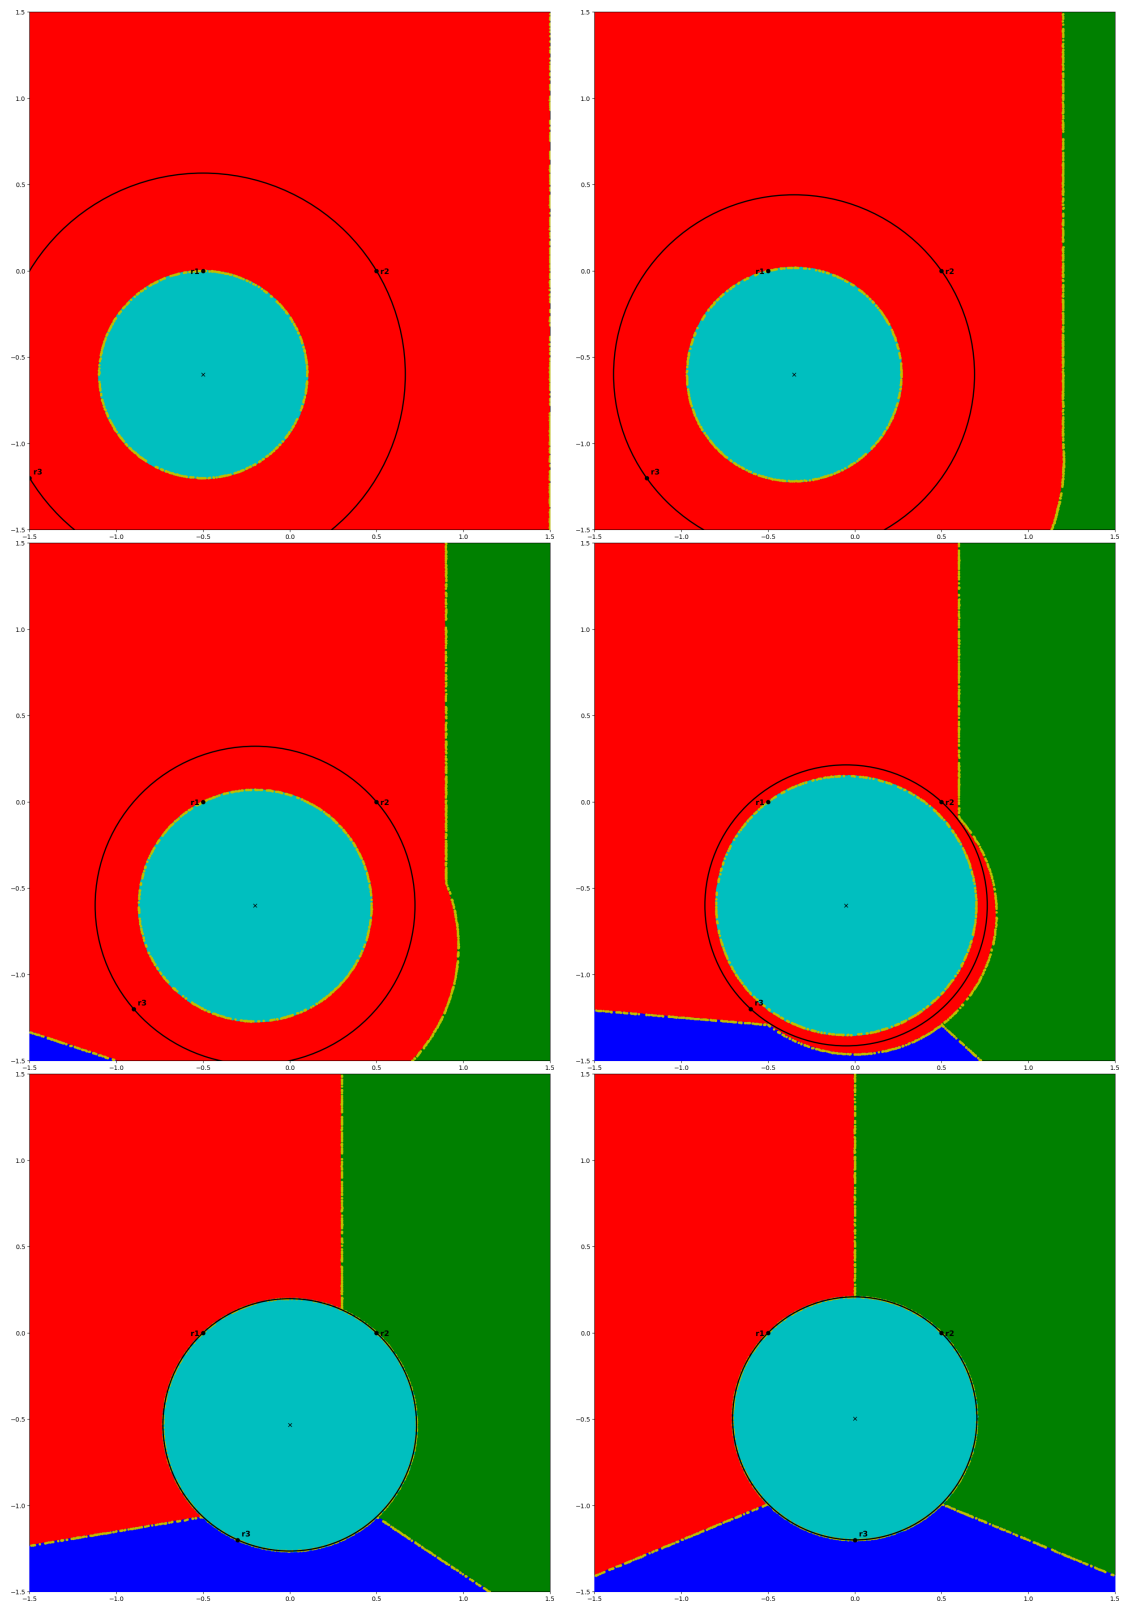


Figure 11

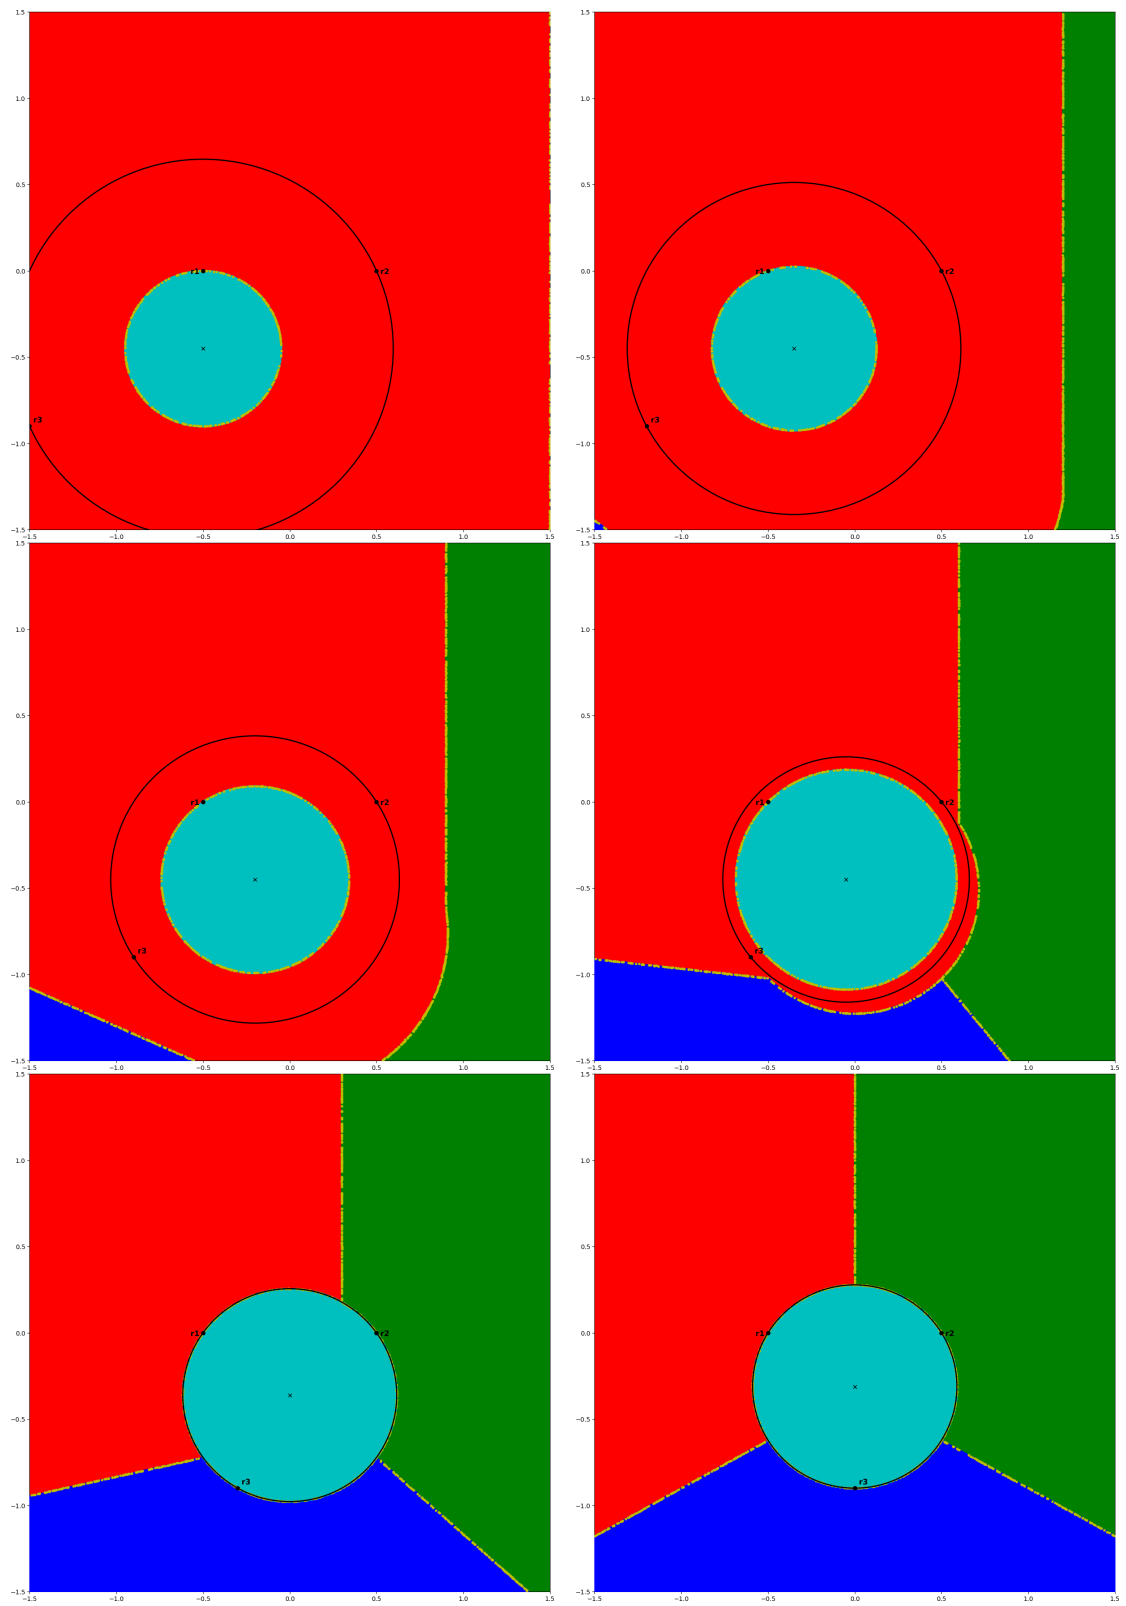


Figure 12




Figure 13

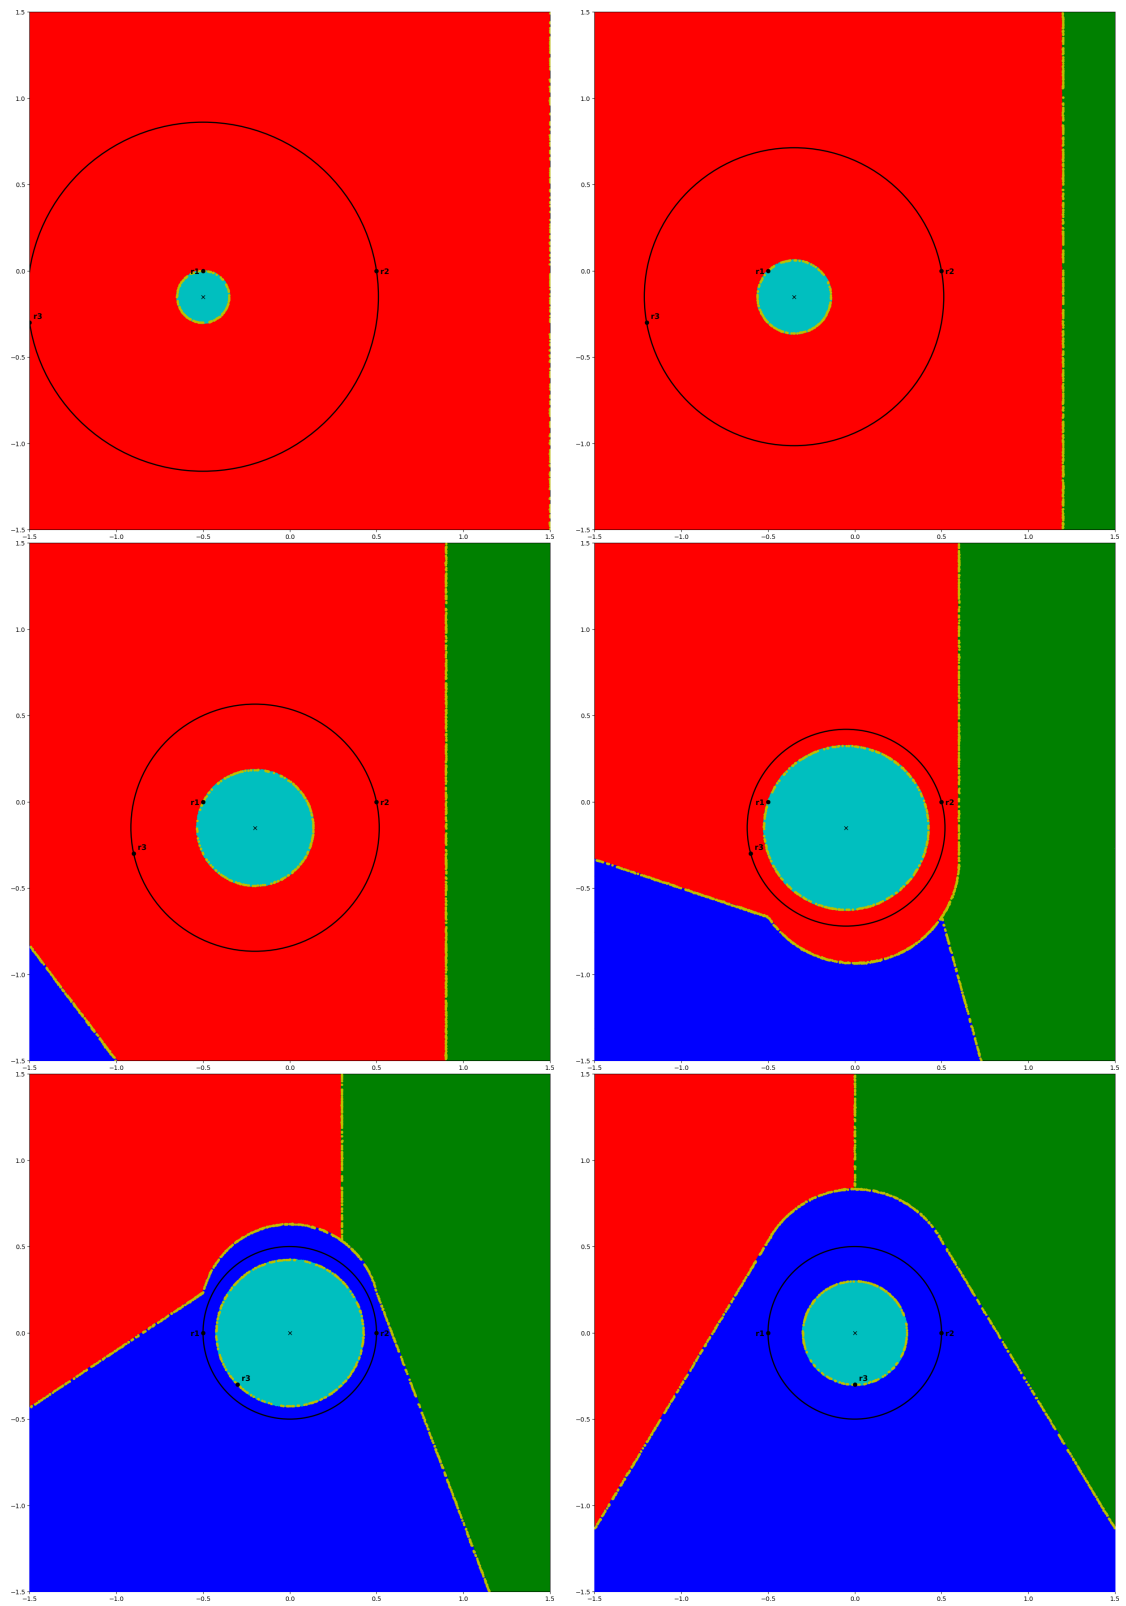


Figure 14




Figure 15