



HAL
open science

Stand Up Indulgent Rendezvous

Quentin Bramas, Anissa Lamani, Sébastien Tixeuil

► **To cite this version:**

Quentin Bramas, Anissa Lamani, Sébastien Tixeuil. Stand Up Indulgent Rendezvous. 22nd International Symposium Stabilization, Safety, and Security of Distributed Systems (SSS 2020), Nov 2020, Austin, TX, United States. pp.45-59, 10.1007/978-3-030-64348-5_4 . hal-02961247

HAL Id: hal-02961247

<https://hal.science/hal-02961247v1>

Submitted on 8 Oct 2020

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.

Stand Up Indulgent Rendezvous^{*}

Quentin Bramas¹, Anissa Lamani¹, and Sébastien Tixeuil²

¹ ICUBE, Strasbourg University, CNRS, France

² Sorbonne University, CNRS, LIP6, France

Abstract. We consider two mobile oblivious robots that evolve in a continuous Euclidean space. We require the two robots to solve the rendezvous problem (meeting in finite time at the same location, not known beforehand) despite the possibility that one of those robots crashes unpredictably. The rendezvous is stand up indulgent in the sense that when a crash occurs, the correct robot must still meet the crashed robot on its last position.

We characterize the system assumptions that enable problem solvability, and present a series of algorithms that solve the problem for the possible cases.

1 Introduction

The study of swarm robotics in Distributed Computing has focused on the computational power of a set of autonomous robots evolving in a bidimensional Euclidean space. In this setting, a robot is modeled as a point in a two dimensional plane and has its own coordinate system and unit distance. Robots are usually assumed to be very weak : they are *(i)* anonymous (they can not be distinguished), *(ii)* uniform (they execute the same algorithm) and, *(iii)* oblivious (they cannot remember past actions). Robots operate in cycles that comprise three phases: Look, Compute and Move. During the first phase (Look), a robot takes a snapshot to see the position of the other robots. During the second phase (Compute), a robot decides to move or stay idle. In the case in which it decides to move, it computes a target destination. In the last phase (Move), a robot moves to the computed destination (if any). Depending on how robots are activated and the synchronization level, three models have been introduced: Fully synchronous model (FFSYNC) in which robots are activated simultaneously and execute cycles synchronously. Semi synchronous model (SSYNC) in which a subset of robots are activated simultaneously. The activated robots execute a cycle synchronously. The asynchronous model (ASYNC) in which there is no global clock. The duration of each phase is finite but unbounded. That is, one robot can decide to move according to an outdated view.

Among the various problems considered under such weak assumptions there is the gathering problem which is one of the benchmarking tasks in mobile robot networks. The gathering task consists in all robots reaching a single point, not known beforehand, in finite time. The particular case of gathering two robots is called *rendezvous* in the literature. In this paper, we consider the Stand Up Indulgent Rendezvous (SUIR) problem: in the case one of the two robots crashes, they still have to gather (obviously, at the position of the crashed robot); if no robot crashes, the robots are expected to gather in finite time. The SUIR problem is at least as difficult as the rendezvous problem, so classical impossibility results still apply.

Related works. A foundational result [14] shows that when robots operate in a fully synchronous manner, the rendezvous can be solved deterministically, while if robots are allowed to wait for a while (this is the case *e.g.* in the SSYNC model), the problem becomes impossible without additional assumptions. Such additional assumptions include the robots executing a probabilistic protocol [8, 9] (but the rendezvous only occurs probabilistically), the robots sharing a common $x-y$ coordinate system [14] or an approximation of a common coordinate system [12], or the robots being endowed with persistent memory [7, 11, 14, 15, 13]. Recent work considered the minimum amount of persistent memory that is necessary to solve rendezvous [7, 11, 13, 15]. It turns out that

^{*} This work was partially funded by the ANR project SAPPORO, ref. 2019-CE25-0005-1.

exactly one bit of persistent memory is necessary and sufficient [11] even when robots operate asynchronously and are disoriented.

When robots can crash unexpectedly, two variants of the gathering problem can be defined [1]: *weak gathering* requires correct robots to gather, regardless of the position of crashed robots; *strong gathering* requires all robots to gather at the same point. Obviously, strong gathering is only feasible if all crashed robots crash at the same location. Early solutions to weak gathering in SSYNC for groups of at least three robots make use of extra hypotheses: (i) starting from a distinct configuration (that is, a configurations where at most one robot occupies a particular position), at most one robot may crash [1], (ii) robots are activated one at a time [8], (iii) robots may exhibit probabilistic behavior [9], (iv) robots share a common chirality (that is, the same notion of handedness) [4], (v) robots agree on a common direction [3]. It turns out that these hypotheses are *not* necessary to solve deterministic weak gathering in SSYNC, when up to $n - 1$ robots may crash [5].

The case of strong gathering mostly yielded impossibility results: with at most a single crash, strong gathering $n \geq 3$ robots deterministically in SSYNC is impossible even if robots are executed one at a time, and probabilistic strong gathering $n \geq 3$ robots is impossible with a fair scheduler [8, 9]. However, probabilistic strong gathering $n \geq 3$ robots becomes possible in SSYNC if the relative speed of the robots is upper bounded by a constant [8, 9].

For the special case of two robots, the strong gathering problem boils down to stand up indulgent rendezvous (SUIR), as presented above. Only few results are known:

1. The algorithm "with probability $\frac{1}{2}$, go to the other robot position" is a probabilistic solution to SUIR in SSYNC [8, 9],
2. The algorithm "go to the other robot position" is a deterministic solution to SUIR in SSYNC when exactly one robot is activated at any time [8, 9].

To this paper, it is unknown whether additional assumptions (*e.g.* a common coordinate system in SSYNC, or FSYNC scheduling) enable deterministic SUIR solvability.

Our contribution. In this paper, we consider the SUIR problem and concentrate at characterizing its deterministic solvability. When robots share a common $x - y$ coordinate system, rendezvous is deterministically solvable in SSYNC [14]: the two robots simply meet at the position of the Northernmost, Easternmost position. Our main impossibility result shows that SUIR cannot be solved deterministically in this setting. Furthermore, it remains impossible if robots have *both* infinite persistent memory (this is a stronger assumption than the classical luminous robot model that permits to solve classical rendezvous in ASYNC [7, 13, 11, 15]) and a common $x - y$ coordinate system. This motivates our focus on the FSYNC setting, where both robots always operate synchronously. Our main positive result is that SUIR is deterministically solvable in FSYNC by oblivious disoriented robots. Our approach is constructive: we first present a simple algorithm for the case the robots share a common coordinate system, and then a more involved solution for the case of disoriented robots.

An interesting byproduct of our work is an oblivious deterministic rendezvous protocol (so, assuming no faults) for the case where robots only agree on a single axis. This complements previous results where robots agree on both axes [14] or approximately agree on both axes [12].

A summary of our results is presented in the following table.

	Rendezvous	SUIR
SSYNC oblivious, disoriented	Impossible [14]	Impossible (Theorem 1)
SSYNC oblivious, common x axis	Possible (Algorithm 1)	Impossible (Theorem 1)
SSYNC oblivious, common $x - y$ axes	Possible [10, 14]	Impossible (Theorem 1)
SSYNC luminous, disoriented	Possible [7, 11, 13–15]	Impossible (Theorem 1)
FSYNC oblivious, disoriented	Possible [2, 6, 14]	Possible (Algorithm 3)

2 Model

We consider two robots, evolving in a Euclidean two-dimensional space. Robots are anonymous and oblivious. The time is discrete, and at each time instant, called round, a subset of the robots is activated and executes a Look-Compute-Move cycle. Each activated robot first Looks at its surroundings to retrieve the position of the other robot in its ego-centered coordinate system. Then, it Computes a target destination, based only on the current position of the other robot. Finally, it moves towards the destination following a straight path.

If the movements are *rigid*, each robot always reaches its destination before the next Look-Compute-Move cycle. Otherwise, movements are *non-rigid*, and an adversary can stop the robot anywhere along the path to its destination, but only after the robot traveled at least a fixed positive distance δ . The value of δ is not known by the robots, and can be arbitrary small, but it does not change during the execution.

In the fully-synchronous model (FSYNC), all correct robots are activated at each round. In the *Semi-synchronous* model (SSYNC), only a non-empty subset of the correct robots may be activated at each round. In this case, we consider only *fair* schedules *i.e.*, schedules where each correct robot is activated infinitely often.

Configurations And Local Views. We consider different settings that impact how a robot retrieves the position of the other robot. Robots might agree on one or both axes of their ego-centered coordinate systems. In other words, robots may have a common North (and possibly a common East direction). They may also have different unit distance.

For the analysis, we assume a global coordinate system Z that is not accessible to the robots. A *configuration* is a set $\{r_1, r_2\}$ containing the positions of both robots in Z . Notice that r_i , $i = 1, 2$, denotes at the same time a robot and its position in \mathbb{R}^2 in the coordinate system Z .

To model the agreement of the robots about their coordinate system, we define the set \mathcal{T} of *indistinguishable transformations*, that modify how robots see the current configuration. If robots agree on both axes and on the unit distance, then \mathcal{T} only contains the identity. If robots do not agree on the unit distance, then \mathcal{T} contains all the (uniform) scaling transformations. If robots do not agree on the x -axis, then \mathcal{T} also contains the reflection along the y -axis. If robots does not agree on any axis, then \mathcal{T} also contains all the rotations. Finally, \mathcal{T} is closed by composition.

We say robots are *disoriented* if robots do not agree on any axis, nor on a common unit distance *i.e.*, \mathcal{T} contains the rotations, scaling, reflection, and their compositions.

In a configuration $\{r_1, r_2\}$ the *local view* V_1 of robot r_1 is obtained by translating the global configuration by $-r_1$ (so that r_1 is seen at position $(0, 0)$ and r_2 is at position $r_2 - r_1$) from which we apply the transformation function h_1 of r_1 . Formally, $V_1 = \{(0, 0), h_1(r_2 - r_1)\}$. Similarly the local view of r_2 is $V_2 = \{(0, 0), h_2(r_1 - r_2)\}$, where $h_2 \in \mathcal{T}$ corresponds to the transformation function of robot r_2 . Notice that the transformation function of a robot r is chosen by an adversary but it does not change over time.

A configuration C is said to be *distinct* if $|C| = 2$.

Configurations And Local Views in One Dimension. The evolving space of the robots can be naturally restricted to a one-dimensional space *i.e.*, robots that evolve on a line. In this case, robots in a configuration C correspond to points in \mathbb{R} instead of \mathbb{R}^2 . Similarly, the set of transformation functions \mathcal{T} contains scaling if robots do not agree on the unit distance, and contains the reflection (or equivalently the π -rotation) if robots do not agree on the orientation of the single axis (*i.e.*, are disoriented).

Algorithms And Executions. An algorithm A is a function mapping local views to destinations. The local view of a robot r is centered and transformed by a function h_r , and when activated, algorithm A outputs r 's destination d in its local view. So to obtain the destination of a robot in the global coordinate system Z , one should apply the inverse transformation h_r^{-1} *i.e.*, the global destination is $r + h_r^{-1}(d)$.

When a non-empty subset of robots S executes an algorithm A in a given configuration C , the obtained configuration C' is the smallest set satisfying³:

$$\begin{aligned} \forall r \in C \setminus S &\Rightarrow r \in C' \\ \forall r \in C \cap S &\Rightarrow r + h_r^{-1}(A(\{(0,0), h_r(r' - r)\})) \in C', \text{ with } r' \in C \setminus \{r\}. \end{aligned}$$

In this case, we write $C \xrightarrow{A} C'$, and say C' is obtained from C by applying A . We say a robot *crashes* at time t if it is not activated at time t , and never activated after time t *i.e.*, a crashed robot stops executing its algorithm and remains at the same position.

An *execution* of algorithm A is an infinite sequence of configurations C_0, C_1, \dots such that $C_i \xrightarrow{A} C_{i+1}$ for all $i \in \mathbb{N}$. We say an execution contains one crashed robot, if one robot crashes at some round t .

The Stand Up Indulgent Rendezvous Problem. An algorithm solves the Stand Up Indulgent Rendezvous (SUIR) problem if, for any initial configuration C_0 and for any execution C_0, C_1, \dots with up to one crashed robot, there exists a round t and a point p such that $C_{t'} = \{p\}$ for all $t' \geq t$.

Informally, if one robot crashes, the correct robot goes to the crashed robot; if no robot crashes, both robots gather in a finite number of rounds.

Since we allow arbitrary initial configuration and the robots are oblivious, we can consider without loss of generality that the crash, if any, occurs at the start of the execution.

3 Impossibility Results

In this section we prove that the SUIR problem is not solvable in SSYNC, even if robots share a full coordinate system (the transformation function is the identity), and have access to infinite persistent memory that is readable by the other robot. In the literature [7], the persistent memory aspect is called the Full-light model with an infinite number of colors. We now prove that having such capabilities does not help solving the problem. The next lemma is very simple but is a key point to prove our main impossibility result.

Lemma 1. *Consider the SSYNC model, with rigid movements, robots endowed with full-lights with infinitely many colors, and a common coordinate system. Assuming algorithm A solves the SUIR problem, then, in every execution suffix starting from a distinct configuration where only robot r is activated (e.g. because the other robot has crashed), there must exist a configuration where algorithm A commands that r moves to the other robot's position.*

Proof. Any move of robot r that does not go to the other robot location does not yield gathering. If this repeats infinitely, no rendezvous is achieved. \square

³ This definition works when $|C| = 2$ but can be easily generalized to larger configurations

Theorem 1. *The SUIR problem is not solvable in SSYNC, even with rigid movements, robots endowed with full-lights with infinitely many colors, and sharing a common coordinate system.*

Proof. Assume for the purpose of contradiction that such an algorithm exists. Let r be one of the robots, and r' be the other robot. We construct a fair infinite execution where rendezvous is never achieved. At some round t , we either activate only r , only r' , or both, depending on what the (deterministic) output of the algorithm in the current configuration is:

- If r is dictated to stay idle: activate only r
- If r is dictated to move to $p \neq r'$: activate only r
- If r is dictated to move to r' , and r' is dictated to move: activate both robots.
- Otherwise (r' is dictated to stay idle): activate only r'

We now show that the execution is fair. Suppose for the purpose of contradiction that the execution is unfair, so there exists a round t after which only r is executed, or only r' is executed. In the first case, it implies there exists an execution suffix where r is never dictated to move to the other robot, which contradicts Lemma 1. Now, if only r' is activated after some round t , then there exists a suffix where r' is always dictated to stay idle, which also contradicts Lemma 1.

The schedule we choose guarantees the following. When only r is activated, rendezvous is not achieved as r is not moving to r' . When only r' is activated, rendezvous is not achieved as r' is idle. If both robots are activated rendezvous is not achieved as r is moving to r' while r' is moving.

Overall, there exists an infinite fair execution where robots never meet, a contradiction with the initial assumption that the algorithm solves SUIR. \square

4 Reduction To One-dimensional Space

In this section, we prove that having an algorithm solving the SUIR problem in a one-dimensional space implies the existence of an algorithm solving the same problem in a two-dimensional space. This theorem is important as our algorithms are defined on the one-dimensional space. However, since we do not prove the converse, the impossibility result we present in the previous Section works in the most general settings, assuming a two-dimensional space. Indeed, we present after the theorem, an example of a problem (the fault-free rendezvous with one common full axis) that is solvable in a two dimensional space, but that cannot be reduced to the one-dimensional space. Despite the results being intuitive, the formal proof is not trivial.

Theorem 2. *Suppose there exists an algorithm solving the SUIR problem where robots are restricted to a one-dimensional space. Then, there exists an algorithm solving the SUIR problem in a two-dimensional space.*

Proof. Let A_1 be an algorithm solving the SUIR problem where robots are restricted to a one-dimensional space. We provide a constructive proof of the theorem by giving a new algorithm A_2 executed by robots in the two-dimensional space.

First, for a configuration $C = \{r_{min}, r_{max}\}$, with $r_{min} < r_{max}$ (using the lexicographical order on their coordinates), we define the function \mathbf{v} as follows:

$$\mathbf{v}(\{r_{min}, r_{max}\}) = \frac{r_{max} - r_{min}}{\|r_{max} - r_{min}\|}$$

Let r_1 and r_2 denote the two robots, having transformation functions h_1 and h_2 , respectively. Let $V_i = \{(0,0), h_i(r_j - r_i)\}$, $i = 1, 2$, $j = 3 - i$, be the local view of robot r_i . Each robot r_i can compute its own orientation vector $v_i = \mathbf{v}(V_i)$ of the line joining the two robots. Notice that, if robots remains on the same line, then v_i remains invariant during the whole execution as long as robots do not gather.

We define algorithm A_2 , executed by robots in the two-dimensional space as follow. First, if the local view of a robot r_i is $\{(0,0)\}$, then A_2 outputs $(0,0)$.

Otherwise, r_i can map its local view V_i in a one-dimensional space to obtain $\overline{V}_i = \{0, b_i\}$ with b_i such that $h_i(r_j - r_i) = b_i v_i$, and execute A_1 on \overline{V}_i . The obtained destination $\overline{p} = A_1(\overline{V}_i)$, is then converted back to the two-dimensional space to obtain the destination $p = \overline{p}v_i$. Doing so, the robots remain on the same line, and v_i remains invariant while the robots are not gathered (when they are gathered, both algorithms stop).

Let $E = C_0, C_1, \dots$ be an arbitrary execution of A_2 . We want to construct from E an execution \overline{E} of A_1 such that the rendezvous is achieved in \overline{E} if and only if the rendezvous is achieved in E .

Recall that we analyze each configuration C_i , $i \in \mathbb{N}$, using Z , the global coordinate system we use for the analysis. Let $v = \mathbf{v}(C_0)$. Again, since robots remain on the same line, then $v = \mathbf{v}(C_i)$ for any C_i while robots are not yet gathered.

Let O be any point of the line L joining the two robots. We define as follow the bijection m mapping points of L (in Z), to the global one-dimensional coordinate system (O, v) :

$$\forall a \in \mathbb{R}, m(O + av) = a$$

We can extend the function m to configurations as follows:

$$m(C) = \{m(r) \mid r \in C\}.$$

Now, let $\overline{E} = \overline{C}_0, \overline{C}_1, \dots$ be the execution of A_1 , in (O, v) , of two robots having transformation function \overline{h}_1 and \overline{h}_2 respectively, with $\overline{h}_i(a) = b$ if and only if $h_i(av) = bv_i$.

We now show that, if $C \xrightarrow{A_2} C'$ then $m(C) \xrightarrow{A_1} m(C')$. To do so we show that the result of executing A_1 on $m(C)$ coincides with $m(C')$. Let $C = \{O + a_1v, O + a_2v\}$, i be an activated robot and j be the other robot. On one hand, we have $m(C) = \{a_1, a_2\}$ and, by construction, the view \overline{V}_i of robot i in $m(C)$ is $\{0, b_i\}$ with $b_i = \overline{h}_i(a_j - a_i)$, so that the global destination of r_i in $m(C)$ is then $a_i + \overline{h}_i^{-1}(\overline{p})$ (with $\overline{p} = A_1(\overline{V}_i)$). On the other hand, the view V_i of robot i is $\{(0, 0), h_i((a_j - a_i)v)\} = \{(0, 0), b_i v_i\}$, so that the global destination of r_i in Z is $O + a_i v + h_i^{-1}(\overline{p}v_i) = O + a_i v + \overline{h}_i^{-1}(\overline{p})v$. Since $m(O + a_i v + \overline{h}_i^{-1}(\overline{p})v) = a_i + \overline{h}_i^{-1}(\overline{p})$, we obtain that $m(C) \xrightarrow{A_1} m(C')$ (assuming the same robots are activated in C and in $m(C)$).

Hence, $\overline{C}_i = m(C_i)$ for all $i \in \mathbb{N}$. Since A_1 solves the SUIR problem, there exists a point $p \in \mathbb{R}$ and a round t such that, for all $t' \geq t$, $m(C_{t'}) = \{p\}$. This implies that $C_{t'} = \{O + pv\}$, so that A_2 solves the SUIR problem. \square

Rendezvous without faults with one full axis. Now, we show that the converse of Theorem 2 is not true in the fault-free model. This observation justifies that, for the results to be more general, we defined our model and gave our impossibility results for the two-dimensional space.

We present an algorithm solving the (fault-free) rendezvous problem in a two-dimensional space, assuming robots agree on one full axis (that is, they agree on the direction and the orientation of the axis). Under this assumption, it is possible that the robots do not agree on the orientation of the line joining them, so that assuming the converse of Theorem 2 would imply the existence of an algorithm in the one-dimensional space with disoriented robots (which does not exist, using a similar proof as the one given in [14]).

The idea is that, if the configuration is symmetric (robots may have the same view), then robots move to the point that forms, with the two robots, an equilateral triangle. Since two such points exist, the robots choose the northernmost one (the robots agree on the y -axis, which provides a common North). Otherwise, the configuration is not symmetric, and there is a unique northernmost robot r . This robot does not move and the other robot moves to r .

Algorithm 1: Fault-free rendezvous. Robots agree on one full axis, may not have common unit distance, and movements are non-rigid

Data: r : robot executing the algorithm

Let $\{(0, 0), (x, y)\}$ be r 's local view.

if $y = 0$ **then**

1 | move to the point $\left(x/2, \left|\frac{x\sqrt{3}}{2}\right|\right)$

else if $y > 0$ **then**

2 | move to the other robot, at (x, y) .

Theorem 3. *Algorithm 1 solves the (fault-free) rendezvous problem with non-rigid movements and robots having only one common full axis, and different unit distance.*

Proof. If the configuration is not symmetric, then, after executing Line 2, either the rendezvous is achieved, or the moving robots remains on the same line joining the robots, so the obtained configuration is still asymmetric, and the same robot is dictated to move towards the same destination, so it reaches it in a finite number of rounds.

Consider now that the initial configuration C is symmetric. If both robots reaches their destination, the rendezvous is completed in one round. If robots are stopped before reaching their destinations, two cases can occur. Either they are stopped after traveling different distances, or they are stopped at the same y -coordinate. In the former case, the obtained configuration is asymmetric and we retrieve the first case. In the latter, the configuration remains symmetric, but the distance between the two robots decreases by at least $\frac{2\delta}{\sqrt{2}}$. Hence, at each round either robots complete the rendezvous, reach an asymmetric configuration, or come closer by a fixed distance. Since the latter case cannot occurs infinitely, one of the other case occurs at least once, and the rendezvous is completed in a finite amount of rounds. \square

5 SUIR Algorithm for FSYNC Robots with a Common Coordinate System

Since it is impossible to solve the SUIR problem in SSYNC, we now concentrate on FSYNC. We first consider a strong model, assuming robots agree on both axis, have a common unit distance, and assuming movements are rigid, before relaxing all hypotheses in Section 6. By Theorem 2, it is sufficient to give an algorithm for a one-dimensional space. Figure 1 illustrates the two possible configurations: either the distance between the two robots is greater than one (the common unit distance), or at most one. In the former case both robots move to the middle. In the latter, we can dictate the right robot to move to the left one, and the left robot to move one unit distance to the right of the other robot. Recall that we can distinguish the left and the right robot on the line because we assume the robots agree on both axis in the two-dimensional space.

Algorithm 2: Rendezvous with rigid movements, and a common coordinate system

Data: r : robot executing the algorithm
 d : the distance between the two robots
if $d > 1$ **then** move to the middle
else
 if r is the left robot **then**
 | move to the point at distance one at the right of the other robot
 else
 | move to the other robot




Fig. 1. The two possible configurations, depending the distance between the two robots

Theorem 4. *Algorithm 2 solves the SUIR problem in FSYNC with rigid movement, and robots agreeing on both axes and unit distance.*

Proof. First we consider the case where a robot crashes. If the left robot crashes, the right robot halves its distance with the other robot each round until its distance is at most one. Then, the right robot move to the other robot and the rendezvous is achieved.

If the right robot crashes, the left robot halves its distance with the other robot each round until its distance is at most one unit. Then, the left robot moves to the right of the other robot, at distance one. It then move to the other robot and the rendezvous is achieved.

Now assume no robot crashes. If the configuration is such that the distance d between the two robots is greater than one, then both robots move to the middle at the same time and the rendezvous is achieved. Otherwise, if the distance d is at most one, after one round the robots are at distance $1 + d > 1$, so that after one more round, the rendezvous is achieved. \square

6 SUIR Algorithm for Disoriented Robots in FSYNC

In this section we present Algorithm 3, which works with disoriented robots (robots do not agree on any axis, nor on the unit distance), and non-rigid movements. The algorithm is defined on the line. Each robot sees the line oriented in some way, but robots might not agree on the orientation of the line. However, since the orientation of the line is deduced from the robot own coordinate system, it does not change over time.

Algorithm 3: SUIR Algorithm for disoriented robots

Let d be the distance to the other robot

Let $i \in \mathbb{Z}$ such that $d \in [2^{-i}, 2^{1-i})$

if $i \equiv 0 \pmod{2}$ **then** move to the middle

if $i \equiv 1 \pmod{4}$ **then** *left* \rightarrow move to middle ; *right* \rightarrow move to other

if $i \equiv 3 \pmod{4}$ **then** *left* \rightarrow move to other ; *right* \rightarrow move to middle




Fig. 2. The four possible configurations, depending the distance between the two robots. We have split the case $i \equiv 0 \pmod{2}$ into two lines to help the reader.

The different moves of a robot r depend on whether a r sees itself on the left or the right of the other robot, and on its *level*. The level of robot at distance d from the other robot (according its own coordinate system, hence its own unit distance), is the integer $i \in \mathbb{Z}$ such that $d \in [2^{-i}, 2^{1-i})$. Figure 2 summarizes the eight possible views of a robot r , and the corresponding movements. Each line represents the congruence of the level of the robot modulo four, and on each line, we see the movement of the robot whether it sees itself on the right or on the left of the other robot. A given figure does not necessarily imply that both robots will actually perform the corresponding movement at the same time (since they may *not* have the same view).

For instance, if a robot r_1 has a level i_1 congruent to 1 modulo 4 and sees itself on the right, while the other robot r_2 has a level i_2 congruent to 2 modulo 4, and also sees itself on the right, then r_1 moves to the other robot position, and r_2 moves to the middle. Assuming both robots reach their destinations, then the distance between them is divided by two (regardless of the coordinate system) so their levels increase by one, and they both see the other robot on the other side, so each robot now sees the other robot on its left.

Let C be any configuration and d is the distance (in the global coordinate system Z) between the two robots. Let x , resp. y , be the distance, in Z , traveled by the left robot, resp. the right robot. Since the robots move toward each other, after executing one round, the distance between the robot becomes $f(d, x, y) = |d - x - y|$.

Lemma 2. *If at least one robot is dictated to move to the middle, then we have $f(d, x, y) \leq d - \min(\delta, d/2)$.*

Proof. For any fixed d , using the symmetry of f (with respect to the second and third argument), we have $f(d, x, y) = g_d(x + y)$ with $g_d : w \mapsto |d - w|$. We know that the distance traveled by the robots is either 0 (if one robot crashes), or at least $\min(d/2, \delta)$, but we cannot have $x = y = 0$. Also, since at least one robot moves to the middle, we have either (i) $x \leq d/2$ and $y \leq d$, or (ii) $x \leq d$ and $y \leq d/2$. Hence, the sum $x + y$ is in the interval $[\min(d/2, \delta), 3d/2]$.

As a convex function, the maximum of g_d is reached at the boundary of its domain

$$\begin{aligned} f(d, x, y) = g_d(x + y) &\leq \max(g_d(3d/2), g_d(\min(d/2, \delta))) \\ &\leq \max(d/2, d - \min(d/2, \delta)) = d - \min(d/2, \delta) \end{aligned}$$

□

The next lemma is a direct consequence of Lemma 2.

Lemma 3. *From a configuration where robots are at distance d (in Z), then, after two rounds, the distance between the robots decreases by at least $\min(\delta, d/2)$.*

Proof. From Lemma 2, we know that if at least one robot is dictated to move to the middle, then, after one round, the distance between the two robots decreases by at least $\min(\delta, d/2)$. Otherwise, we know that both robots were dictated to move to the other robot location. Hence, the distance between the robots is either at most $d - \min(\delta, d/2)$, or is greater than $d - \min(\delta, d/2)$. In the former case, the lemma is proved. In the latter case the order of the robots changes (a left robot becomes a right robot, and vice and versa). This happens regardless of their coordinates system (maybe both robots view themselves on the right, then they both view themselves at the left). Also, the level of each robot is either the same, or is increased by one. In all cases, both robots are dictated to move to the middle. In more details, a right robot with level $i \equiv 1 \pmod{4}$ becomes a left robot with either the same level of level $i + 1 \equiv 2 \pmod{4}$. In both cases, in the next round, the robot is dictated to move to the middle. A left robot with level $i \equiv 3 \pmod{4}$ becomes a right robot with either the same level of level $i + 1 \equiv 0 \pmod{4}$. In both cases, in the next round, the robot is dictated to move to the middle.

So that after one more round, the distance decreases by at least $\min(\delta, d/2)$. □

Lemma 4. *Assuming rigid movement and one robot crash, Algorithm 3 solves the SUIR problem.*

Proof. Let i be the level of the correct robot r . Assume the other robot crashes. Robot r either sees itself on the right or on the left of the other robot.

If r sees itself on the right, then depending on its level, either r moves to the middle, or move to the other robot. In the former case, the level of r increases by one and r continues to see itself on the right. In the latter case, the rendezvous is achieved in one round. After at most three rounds, the level of r is congruent to 1 modulo 4 so that after at most four rounds the rendezvous is achieved.

Similarly, if r sees itself on the left, then after at most four rounds, r level is congruent to 3 modulo 4 and the rendezvous is achieved. □

If, at round t , one robot sees itself on the right, and the other sees itself on the left, then they agree on the orientation of the line at time t . Since, for each robot, the orientation of the line does not change, then they agree on it during the whole execution (except when they are gathered, as the line is not defined in that case).

Similarly, if at some round, both robots see themselves at the right (resp. at the left), then their orientations of the line are opposite, and remain opposite during the whole execution (again, until they gather). Hence we have the following remark.

Remark 1. Consider two disoriented robots moving on the line L joining them and executing Algorithm 3. Then, either they have a common orientation of L during the whole execution (while they are not gathered), or they have opposite orientations of L during the whole execution (while they are not gathered).

Lemma 5. *Assuming rigid movements, no crash, and robots having a **common orientation** of the line joining them, then, Algorithm 3 solves the SUIR problem.*

Proof. Since the robots have a common orientation, we know there is one robot that sees itself on the right and one robot that sees itself on the left. Of course the robots are not aware of this, but we saw in the previous remark that a common orientation is preserved during the whole execution (while robots are not gathered).

Let $(i, j) \in \mathbb{Z}^2$ denote a configuration where the robot on the left is at level i , and the robot on the right is at level j , and we write $(i, j) \equiv (k, l) \pmod{4}$ if and only if $i \equiv k \pmod{4}$ and $j \equiv l \pmod{4}$.

To prove the lemma we want to show that for any configuration $(i, j) \in \mathbb{Z}^2$, the robots achieve rendezvous. Take an arbitrary configuration $(i, j) \in \mathbb{Z}^2$. We consider all 16 cases:

1. **if** $(i, j) \equiv (0, 0) \pmod{4}$: rendezvous is achieved in one round.
2. **if** $(i, j) \equiv (0, 1) \pmod{4}$: we reach configuration $(j + 1, i + 1) \equiv (2, 1) \pmod{4}$
3. **if** $(i, j) \equiv (0, 2) \pmod{4}$: rendezvous is achieved in one round.
4. **if** $(i, j) \equiv (0, 3) \pmod{4}$: rendezvous is achieved in one round.
5. **if** $(i, j) \equiv (1, 0) \pmod{4}$: rendezvous is achieved in one round.
6. **if** $(i, j) \equiv (1, 1) \pmod{4}$: we reach configuration $(j + 1, i + 1) \equiv (2, 2) \pmod{4}$
7. **if** $(i, j) \equiv (1, 2) \pmod{4}$: rendezvous is achieved in one round.
8. **if** $(i, j) \equiv (1, 3) \pmod{4}$: rendezvous is achieved in one round.
9. **if** $(i, j) \equiv (2, 0) \pmod{4}$: rendezvous is achieved in one round.
10. **if** $(i, j) \equiv (2, 1) \pmod{4}$: we reach configuration $(j + 1, i + 1) \equiv (2, 3) \pmod{4}$
11. **if** $(i, j) \equiv (2, 2) \pmod{4}$: rendezvous is achieved in one round.
12. **if** $(i, j) \equiv (2, 3) \pmod{4}$: rendezvous is achieved in one round.
13. **if** $(i, j) \equiv (3, 0) \pmod{4}$: we reach configuration $(j + 1, i + 1) \equiv (1, 0) \pmod{4}$
14. **if** $(i, j) \equiv (3, 1) \pmod{4}$: we reach configuration $(j, i) \equiv (1, 3) \pmod{4}$
15. **if** $(i, j) \equiv (3, 2) \pmod{4}$: we reach configuration $(j + 1, i + 1) \equiv (3, 0) \pmod{4}$
16. **if** $(i, j) \equiv (3, 3) \pmod{4}$: we reach configuration $(j + 1, i + 1) \equiv (0, 0) \pmod{4}$

In any case, rendezvous is achieved after at most three rounds. □

Lemma 6. *Assuming rigid movement, no crash, and robots having **opposite orientations** of the line joining them, then, Algorithm 3 solves the SUIR problem.*

Proof. Since the robots have opposite orientations, we know they either both see themselves on the right or they both see themselves on the left. Of course the robots are not aware of this, but we saw in the previous remark that the opposite orientations are preserved during the whole execution (while robots are not gathered).

In this proof, $R\{i, j\}$ denotes a configuration where both robots see themselves on the right and one of them has level i and the other as level j . Here, the order between i and j does not matter (hence the set notation). Similarly $L\{i, j\}$ denotes a configuration where both robots see themselves on the left, and one of them has level i and the other as level j .

Here, assuming without loss of generality that $i \leq j \pmod{4}$, we write $R\{i, j\} \equiv (k, l) \pmod{4}$, resp. $L\{i, j\} \equiv (k, l) \pmod{4}$, if and only if, $i \equiv k \pmod{4}$ and $j \equiv l \pmod{4}$.

To prove the lemma, we want to show that for any configuration $R\{i, j\}$ or $L\{i, j\}$, the robots achieve rendezvous. Take an arbitrary configuration $(i, j) \in \mathbb{Z}^2$. We consider all 20 cases:

1. **if** $L\{i, j\} \equiv (0, 0) \pmod 4$: rendezvous is achieved in one round.
2. **if** $L\{i, j\} \equiv (0, 1) \pmod 4$: rendezvous is achieved in one round.
3. **if** $L\{i, j\} \equiv (0, 2) \pmod 4$: rendezvous is achieved in one round.
4. **if** $L\{i, j\} \equiv (0, 3) \pmod 4$: we reach configuration $R\{i+1, j+1\} \equiv (0, 1) \pmod 4$.
5. **if** $L\{i, j\} \equiv (1, 1) \pmod 4$: rendezvous is achieved in one round.
6. **if** $L\{i, j\} \equiv (1, 2) \pmod 4$: rendezvous is achieved in one round.
7. **if** $L\{i, j\} \equiv (1, 3) \pmod 4$: we reach configuration $R\{i+1, j+1\} \equiv (0, 2) \pmod 4$.
8. **if** $L\{i, j\} \equiv (2, 2) \pmod 4$: rendezvous is achieved in one round.
9. **if** $L\{i, j\} \equiv (2, 3) \pmod 4$: we reach configuration $R\{i+1, j+1\} \equiv (0, 3) \pmod 4$.
10. **if** $L\{i, j\} \equiv (3, 3) \pmod 4$: we reach configuration $R\{i, j\} \equiv (3, 3) \pmod 4$.
11. **if** $R\{i, j\} \equiv (0, 0) \pmod 4$: rendezvous is achieved in one round.
12. **if** $R\{i, j\} \equiv (0, 1) \pmod 4$: we reach configuration $L\{i+1, j+1\} \equiv (1, 2) \pmod 4$.
13. **if** $R\{i, j\} \equiv (0, 2) \pmod 4$: rendezvous is achieved in one round.
14. **if** $R\{i, j\} \equiv (0, 3) \pmod 4$: rendezvous is achieved in one round.
15. **if** $R\{i, j\} \equiv (1, 1) \pmod 4$: we reach configuration $L\{i, j\} \equiv (1, 1) \pmod 4$.
16. **if** $R\{i, j\} \equiv (1, 2) \pmod 4$: we reach configuration $L\{i+1, j+1\} \equiv (2, 3) \pmod 4$.
17. **if** $R\{i, j\} \equiv (1, 3) \pmod 4$: we reach configuration $L\{i+1, j+1\} \equiv (0, 2) \pmod 4$.
18. **if** $R\{i, j\} \equiv (2, 2) \pmod 4$: rendezvous is achieved in one round.
19. **if** $R\{i, j\} \equiv (2, 3) \pmod 4$: rendezvous is achieved in one round.
20. **if** $R\{i, j\} \equiv (3, 3) \pmod 4$: rendezvous is achieved in one round.

In any case, rendezvous is achieved after at most three rounds. □

Theorem 5. *Algorithm 3 solves the SUIR problem with disoriented robots in FSYNC.*

Proof. By Lemma 3, the distance between the two robots decreases by at least $\min(\delta, d/2)$ every two rounds. Hence, eventually, robots are at distance smaller than δ from one another and, from this point in time, movements are rigid. Assume now that movements are rigid. If a robot crashes, then the rendezvous is achieved by using Lemma 4. Otherwise, depending on whether the robots have a common orientation or opposite orientation of the line joining them (see Remark 1), the Theorem follows either by using Lemma 5 or by using Lemma 6. □

7 Concluding remarks

We considered the problem of Stand Up Indulgent Rendezvous (SUIR). Unlike classical rendezvous, the SUIR problem is unsolvable in SSYNC even with the strongest assumptions: robots share a common $x - y$ coordinate system, and have access to infinite persistent memory. We demonstrate that it is nevertheless solvable in FSYNC without *any* additional assumptions. A natural open question is related to the optimality (in time) of our algorithm.

Also, we would like to investigate further the possibility of deterministic strong gathering for $n \geq 3$ robots. It is known that executing a single robot at a time in SSYNC is insufficient [8, 9], but additional hypotheses may make the problem solvable.

References

1. Agmon, N., Peleg, D.: Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM J. Comput.* **36**(1), 56–82 (2006). <https://doi.org/10.1137/050645221>
2. Balabonski, T., Delga, A., Rieg, L., Tixeuil, S., Urbain, X.: Synchronous gathering without multiplicity detection: a certified algorithm. *Theory Comput. Syst.* **63**(2), 200–218 (2019). <https://doi.org/10.1007/s00224-017-9828-z>
3. Bhagat, S., Chaudhuri, S.G., Mukhopadhyaya, K.: Fault-tolerant gathering of asynchronous oblivious mobile robots under one-axis agreement. In: Rahman, M.S., Tomita, E. (eds.) *WALCOM: Algorithms and Computation - 9th International Workshop, WALCOM 2015, Dhaka, Bangladesh, February 26-28, 2015. Proceedings. Lecture Notes in Computer Science*, vol. 8973, pp. 149–160. Springer (2015)

4. Bouzid, Z., Das, S., Tixeuil, S.: Gathering of mobile robots tolerating multiple crash faults. In: Proc. 33rd IEEE Intl. Conf. on Distributed Computing Systems (ICDCS). pp. 337–346 (Jul 2013). <https://doi.org/10.1109/ICDCS.2013.27>
5. Bramas, Q., Tixeuil, S.: Wait-free gathering without chirality. In: International Colloquium on Structural Information and Communication Complexity. pp. 313–327. Springer (2015)
6. Cohen, R., Peleg, D.: Convergence properties of the gravitational algorithm in asynchronous robot systems. *SIAM J. Comput.* **34**(6), 1516–1528 (2005). <https://doi.org/10.1137/S0097539704446475>
7. Das, S., Flocchini, P., Prencipe, G., Santoro, N., Yamashita, M.: Autonomous mobile robots with lights. *Theor. Comput. Sci.* **609**, 171–184 (2016). <https://doi.org/10.1016/j.tcs.2015.09.018>
8. Défago, X., Gradinariu, M., Messika, S., Parvédy, P.R.: Fault-tolerant and self-stabilizing mobile robots gathering. In: Dolev, S. (ed.) Distributed Computing, 20th International Symposium, DISC 2006, Stockholm, Sweden, September 18–20, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4167, pp. 46–60. Springer (2006). https://doi.org/10.1007/11864219_4
9. Défago, X., Potop-Butucaru, M.G., Clément, J., Messika, S., Parvédy, P.R.: Fault and byzantine tolerant self-stabilizing mobile robots gathering - feasibility study -. *CoRR* **abs/1602.05546** (2016), <http://arxiv.org/abs/1602.05546>
10. Flocchini, P., Prencipe, G., Santoro, N., Widmayer, P.: Gathering of asynchronous robots with limited visibility. *Theor. Comput. Sci.* **337**(1-3), 147–168 (2005). <https://doi.org/10.1016/j.tcs.2005.01.001>
11. Heriban, A., Défago, X., Tixeuil, S.: Optimally gathering two robots. In: Proc. 19th Intl. Conf. on Distributed Computing and Networking, ICDCN. pp. 3:1–3:10 (Jan 2018). <https://doi.org/10.1145/3154273.3154323>
12. Izumi, T., Souissi, S., Katayama, Y., Inuzuka, N., Défago, X., Wada, K., Yamashita, M.: The gathering problem for two oblivious robots with unreliable compasses. *SIAM J. Comput.* **41**(1), 26–46 (2012). <https://doi.org/10.1137/100797916>
13. Okumura, T., Wada, K., Katayama, Y.: Brief announcement: Optimal asynchronous rendezvous for mobile robots with lights. In: Proc. 19th Intl. Symp. on Stabilization, Safety, and Security of Distributed Systems (SSS). pp. 484–488 (Nov 2017). https://doi.org/10.1007/978-3-319-69084-1_36
14. Suzuki, I., Yamashita, M.: Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM J. Comput.* **28**(4), 1347–1363 (1999). <https://doi.org/10.1137/S009753979628292X>
15. Viglietta, G.: Rendezvous of two robots with visible bits. In: Proc. 9th Intl. Symp. on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics, (ALGOSENSORS). pp. 291–306 (Sep 2013). https://doi.org/10.1007/978-3-642-45346-5_21