JAIST Repository
https://dspace.jaist.ac.jp/

Title	The cost of probabilistic agreement in oblivious robot networks
Author(s)	Clement, Julien; Defago, Xavier; Potop-Butucaru, Maria Gradinariu; Izumi, Taisuke; Messika, Stephane
Citation	Information Processing Letters, 110(11): 431-438
Issue Date	2010-05-16
Type	Journal Article
Text version	author
URL	http://hdl.handle.net/10119/9046
Rights	NOTICE: This is the author ' s version of a work accepted for publication by Elsevier. Changes resulting from the publishing process, including peer review, editing, corrections, structural formatting and other quality control mechanisms, may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in Julien Clement, Xavier Defago, Maria Gradinariu Potop-Butucaru, Taisuke Izumi, and Stephane Messika, Information Processing Letters, 110(11), 2010, 431-438, http://dx.doi.org/10.1016/j.ipl.2010.04.006
Description	

Japan Advanced Institute of Science and Technology

The Cost of Probabilistic Agreement in Oblivious Robot Networks

Julien Clement** Xavier Défago† Maria Gradinariu Potop-Butucaru*

clement@lri.fr defago@jaist.ac.jp maria.gradinariu@lip6.fr

Taisuke Izumi†† Stephane Messika**

t-izumi@nitech.ac.jp messika@lri.fr

* Université Pierre et Marie Curie–Paris 6; LIP6/CNRS UMR 7606; 104 av. du Président Kennedy; 75016 Paris, France
** Université Paris-Sud–Paris 11; LRI/CNRS UMR 8623; Laboratoire de Recherche en Informatique; 91405 Orsay, France
† JAIST; School of Information Science; Asahidai 1-1, Nomi; Ishikawa 923-1292, Japan
†† Nagoya Institute of Technology; Gokiso-cho Showa; Nagoya; Aichi 466-8555, Japan

Abstract

In this paper, we look at the time complexity of two agreement problems in networks of oblivious mobile robots, namely, at the gathering and scattering problems. Given a set of robots with arbitrary initial locations and no initial agreement on a global coordinate system, *gathering* requires that all robots reach the exact same but not predetermined location. In contrast, *scattering* requires that no two robots share the same location. These two abstractions are fundamental coordination problems in cooperative mobile robotics. Oblivious solutions are appealing for self-stabilization since they are self-stabilizing at no extra cost. As neither gathering nor scattering can be solved deterministically under arbitrary schedulers, probabilistic solutions have been proposed recently.

The contribution of this paper is twofold. First, we propose a detailed time complexity analysis of a modified probabilistic gathering algorithm. Using Markov chains tools and additional assumptions on the environment, we prove that the convergence time of gathering can be reduced from $O(n^2)$ (the best-known bound) to $O(1)$ or $O(\log n \cdot \log(\log n))$, depending on the model of multiplicity detection. Second, using the same technique, we prove that scattering can also be achieved in fault-free systems with the same bounds.

1 Introduction

Many future applications of mobile robotics envision groups of mobile robots self-organizing and cooperating toward the resolution of common objectives. In many cases, such groups of robots are aimed at being deployed in adverse environments, such as in space, in deep sea, or after disasters (natural or not). Thus, a group must be able to self-organize in the absence of any prior infrastructure (e.g., no global positioning), and ensure coordination in spite of the presence of faulty elements among the robots, as well as other unanticipated environmental changes.

Suzuki and Yamashita [8] proposed a formal model to analyze and prove the correctness of agreement problems in robot networks. In their model, robots are represented as points that evolve on a plane devoid of any landmarks. At any given time, a robot can be either idle or active. When a robot becomes active, it observes the locations of the other robots, computes a target position, and moves towards it. The time when a robot becomes active is governed by an activation scheduler (or daemon). Between two activations robots forget their past computations (robots are said to be *oblivious*). Interestingly, any algorithm proved correct in this model is also inherently self-stabilizing.

The *gathering problem*, also known as the *Rendezvous* problem, is fundamental for coordination in oblivious mobile robotics. Briefly, given a set of robots with arbitrary initial locations and no initial agreement on a global coordinate system, gathering requires that all robots, following their algorithm, reach the exact same location. That location, not agreed upon initially, must be reached within a *finite* number of cycles, with all robots remaining there afterward. The dual problem of gathering is the *scattering problem*. Scattering requires that, starting from an arbitrary configuration, eventually no two robots share the same position.

It turns out that neither deterministic gathering nor scattering are possible without additional assumptions. Most of the work done so far in order to circumvent this impossibility focuses on the additional assumptions the system needs. Surprisingly, the use of randomization has however drawn only little attention so far. No formal framework was proposed in order to analyze the correctness and the complexity of probabilistic algorithms designed for robots networks. In a companion paper [1], we investigated some of the fundamental limits of deterministic and probabilistic gathering in the face of a wide range of synchrony and fault assumptions. Probabilistic scattering was analyzed for the first time by Dieudonné and Petit [2]. However, neither work proposed an actual *framework* in which to analyze the complexity of proposed solutions.

In this paper, we advocate the use of Markov chains as a simple and efficient tool to analyze and compare probabilistic strategies in oblivious robot networks. Note that, in oblivious robot networks, computations can depend only on the current view of the robots, i.e., without making any reference to the past. This behavior makes Markov chains an appealing tool for analyzing their correctness and complexity since, by definition, a Markov chain models a system wherein the next configuration depends strictly on the current one. The only difficulty with using Markov chains consists in associating an appropriate Markov chain to each probabilistic strategy. In this work, we focus on the analysis of existing probabilistic strategies for scattering and gathering. We also claim that our analysis can be easily applied to a broad class of probabilistic strategies (e.g., leader election, flocking, constrained scattering, pattern formation).

Contribution. We show that the time complexity of probabilistic gathering in a fault-free environment can be improved from $O(n^2)$ to as little as $O(1)$ rounds[1] when the algorithms rely on additional information related to the environment (e.g., strong multiplicity knowledge). And, even when information on multiplicity is incomplete (weak multiplicity), we still show a convergence of $O(\log n \cdot \log(\log n))$. We also show the exact same bounds for the scattering problem.

Structure. The paper is structured as follows. Section 2 describes the robot network and system model. Section 3 formally defines the gathering and scattering problems. Section 4 presents the analysis framework. We then analyze the convergence of gathering and scattering both under strong multiplicity (Sect. 5) and under weak multiplicity (Sect. 6). Section 7 concludes the paper and discusses some open problems.

2 Model

We now present the system model considered throughout the paper. The model of robots, and most of the definitions, are due to Suzuki and Yamashita [8] and Prencipe [6].

Robot networks. The system consists of a finite set of robots modeled as dimensionless points moving on a two-dimensional plane. The robots have arbitrary initial locations. They are capable of sensing the environment, computing a position, and moving toward a destination. When sensing, a robot can determine the position of other robots relative to its own local coordinate system.

In this paper, robots are said to have *unlimited visibility*, in the sense that they are always able to sense the position of all other robots, regardless of their proximity.

Multiplicity detection. When several robots share the same location, this location is called a *point of multiplicity*.

Robots are said to have *strong multiplicity knowledge* when they are aware of the number of robots located at each point of multiplicity. In contrast, when robots have *weak multiplicity knowledge*, they know which points are points of multiplicity, but are unable to count how many robots are located there.

Section 5 assumes strong multiplicity knowledge, while Section 6 assumes weak multiplicity.

System model. We consider the model first introduced by Suzuki and Yamashita [8], called SYm, in which robots are repeatedly active and inactive. When a robot becomes active, it performs an *atomic computational cycle* composed of the following three actions: observation, computation, and motion.

- *Observation.* An observation returns a snapshot of the positions of all robots.
- *Computation.* Using the observed environment, a robot executes its algorithm to compute a destination.
- *Motion.* The robot moves to this destination (by a non-zero distance but without always reaching it).

The model considers discrete time at irregular intervals. At each time, some subset of the robots become active and complete an entire computation cycle. Robots can be active either simultaneously or sequentially. Two

[1] A round is the shortest fragment of execution where all robots are activated at least once.

robots that are active simultaneously observe the exact same environment (according to their respective coordinate systems).

Moreover, robots are assumed to be *oblivious* (i.e., stateless), in the sense that a robot does not keep any information between two different computational cycles.

The algorithm of robots is expressed with an I/O automaton [3, 4]. The local state of a robot at time t is the state of its input/output variables and the state of its local variables and registers. A network of robots is modeled by the parallel composition of the individual automaton of each robot. A configuration of the system at time t is the union of the local states of the robots in the system at time t. An execution $e = (c_0, \ldots, c_t, \ldots)$ of the system is an infinite sequence of configurations, where c_0 is the initial configuration,[2] and every transition $c_i \to c_{i+1}$ is associated to the execution of a subset of the previously defined computational cycles.

Schedulers. A scheduler decides at each configuration the set of robots allowed to perform their actions. A scheduler is said to be fair if, in an infinite execution, a robot is activated infinitely often. A scheduler is *centralized*, if it ensures that at most one single robot is active at any given time, otherwise it is *distributed*.

3 Gathering and Scattering

A network of robots is in a *legitimate configuration* with respect to the requirements of gathering if all robots in the system share the same position on the plane. Let us denote by $\mathcal{P}_{Gathering}$ this predicate. An algorithm solves the gathering problem in an oblivious system if the following two properties are verified:

- **Convergence.** Any execution of the system starting in an arbitrary configuration reaches in a finite number of cycles a configuration that satisfies $\mathcal{P}_{Gathering}$.
- **Closure.** Any execution starting in a legitimate configuration with respect to $\mathcal{P}_{Gathering}$ contains only legitimate configurations.

Gathering is difficult to achieve in most environments. Therefore, weaker forms of gathering have also been studied. One interesting variant requires robots to *converge asymptotically* toward a single location, rather than reach such a location in finite time. The convergence is however considerably easier to deal with. For instance, with unlimited visibility, convergence can be achieved trivially by having robots move to the barycenter of the network [8].

The scattering problem was first introduced by Suzuki and Yamashita [7]. The problem aims at arranging a set of robots such that eventually no two robots share the same position. Let us denote by $\mathcal{P}_{Scattering}$ this predicate. In the sequel, we rely on the following definition proposed by Dieudonné and Petit [2]:

- **Convergence.** Any execution of the system starting in an arbitrary configuration reaches in a finite number of cycles a configuration that satisfies $\mathcal{P}_{Scattering}$.
- **Closure.** Any execution starting in a legitimate configuration with respect to $\mathcal{P}_{Scattering}$ contains only legitimate configurations.

4 Analysis framework

In this section, we introduce further definitions needed to analyze the convergence time of probabilistic gathering and scattering. A detailed description of these notions can be found in the literature [5].

Random variables We denote by X_n a random variable. For instance, in our case this could represent the number of groups of size x after n cycles of the algorithm. We will study a discrete-time stochastic process, that is, a sequence $\{X_n\}_{n\geq 0}$ of random variables.

In the sequel, we use the following notation:

- $\mathbb{P}[X_n = x]$ is the probability of the event $\{X_n = x\}$.
- $\mathbb{E}[X_n]$ is the expectation of X_n.
- $X : k \mapsto \mathbb{P}[X = k]$ denotes the probability distribution of a random variable X.

[2] Unless stated otherwise, we make no specific assumption regarding the respective positions of robots in initial configurations.

- $\mathbb{P}[A \mid B]$ is the conditional probability that reads: "the probability of A, given B".

Markov chains Markov chains form a specific class of stochastic processes, with the following fundamental property: the probabilistic dependence on the past is only related to the previous state.

Definition 1 *Let $(X_n)_{n\in\mathbb{N}}$ be a discrete time stochastic process with countable state space E. A stochastic process is called a Markov chain if, for all integers $n \geq 0$ and all states $i_0, i_1, \ldots, i_{n-1}, i, j$, we have:*

$$\mathbb{P}[X_{n+1} = j \mid X_n = i, X_{n-1} = i_{n-1}, \ldots, X_0 = i_0] \;=\; \mathbb{P}[X_{n+1} = j \mid X_n = i]$$

In this paper, we advocate that Markov chains constitute a simple verification tool, well-adapted to the analysis of distributed strategies in oblivious robot networks since, in such networks, the next move of a robot depends only on its current position.

Asynchronous rounds and moving distance. An *asynchronous round* is defined as the shortest fragment of an execution in which each process in the system executes its actions at least once. Throughout this paper, we adopt the number of asynchronous rounds as the unit to evaluate the time complexity of algorithms, which is a standard criterion for asynchronous distributed systems. Although the maximal distance that a robot can reach in a single round is fixed by the model, note that the probability of two robots being at moving distance of each other does not depend on the number of robots. So, in order to simplify the discussion for the gathering analysis, we build our argument on the case where all robots are initially within moving distance of each other.

5 The Case of Strong Multiplicity Knowledge

In this section, we assume that robots have *strong multiplicity knowledge*.

To achieve scattering, we want robots located at the same point x_0 to move randomly to different destinations. If we are able to generate a sufficiently large number of "possible destinations", it is more likely that the robots will move to distinct locations. Therefore, to be efficient, it is desirable that the number of these "possible destinations" increases according to the multiplicity of x_0. Since robots have *strong* multiplicity knowledge, they are aware of the number of robots located at each multiplicity point.

In gathering, we require all robots to gather at a single location. Since robots have *strong* multiplicity knowledge, they are able to determine the best location to gather, i.e, the point of maximal multiplicity.

5.1 Probabilistic scattering

Dieudonné and Petit [2] proved that deterministic scattering is impossible in the SYm model without additional assumptions, and proposed an original probabilistic solution based on the use of Voronoi diagrams, described and analysed in Section 6.1.

Definition 2 *Let $\mathcal{P} = \{p_1, p_2, \ldots, p_n\}$ be a set of points in the Cartesian 2-dimensional plane. The Voronoi diagram of $\mathcal{P}$ is a subdivision of the plane into n cells, one for each point in $\mathcal{P}$. The cells have the property that an arbitrary point q belongs to the Voronoi cell of point p_i if and only if, for any other point $p_j \in \mathcal{P}$, $dist(q, p_i) < dist(q, p_j)$ where $dist(p, q)$ is the Euclidean distance between p and q. In particular, the strict inequality means that points located on the boundary of the Voronoi diagram belong to no Voronoi cell.*

In this section, we modify the scattering algorithm of [2] in order to improve the complexity. The idea of our algorithm referred in the following Algorithm 5.1 is as follows. Robots choose uniformly at random a point within $2n^2$ different positions in their Voronoi cell. Note that Algorithm 5.1 uses implicitly the strong multiplicity knowledge in the computation of the set of possible choice positions, i.e. n. If we assume n known or the knowledge of an upper bound of n then Algorithm 5.1 works under a weak multiplicity model.

Theorem 1 *The expected convergence time of Algorithm 5.1 is $O(1)$.*

Proof: Let $[t_0, t_1]$ be the shortest fragment in which each robot is activated at least once. Let $\mathcal{P} = \{p_1, p_2, \cdots, p_w\}$ be the set of points occupied by two or more robots at t_0, where $w = |\mathcal{P}|$. We define the indicator random variable Z_i as follows: $Z_i = 1$ if all robots located at the same point p_i are located on different points at time t_1, and

Algorithm 5.1 Probabilistic Scattering executed by robot r_i with strong multiplicity knowledge.

Compute the Voronoi diagram;
$Cell_i :=$ Voronoi cell where r_i is located;
$Current\_Pos :=$ position where r_i is located;

Let Pos be a set of $2n^2$ positions in $Cell_i$:
Move toward a position in Pos chosen uniformly at random,

$Z_i = 0$ otherwise. Notice that$\{Z_1, Z_2, \cdots\}$ are mutually independent because we use Voronoi diagrams and thus any two robots from different points never reach the same position.

Let m_i be the multiplicity of p_i at t_0. To prove the theorem, we show that the probability $\mathbb{P}[\bigwedge_{i=1}^{w} Z_i = 1]$ is bounded by a constant. Since two robots moving at different time necessarily have different destinations, the worst case scenario is when m_i robots are activated simultaneously during the interval $[t_0, t_1]$. Thus, the probability $\mathbb{P}[Z_i = 1]$ is bounded as follows:

$$\begin{aligned}\mathbb{P}[Z_i = 1] &\geq \left(\frac{2n^2-1}{2n^2}\right)\left(\frac{2n^2-2}{2n^2}\right)\cdots\left(\frac{2n^2-m_i}{2n^2}\right)\\ &\geq \left(1-\frac{m_i}{2n^2}\right)^{m_i} \geq \left(1-\frac{m_i^2}{2n^2}\right)\end{aligned}$$

The random variables being independent, we get:

$$\mathbb{P}[\bigwedge_{i=1}^{w} Z_i = 1] \geq \prod_{i=1}^{w}\left(1-\frac{m_i^2}{2n^2}\right) \geq \prod_{i=1}^{w}\left(1-\frac{m_i}{2n}\right)$$

Therefore, we can write :

$$log(\mathbb{P}[\bigwedge_{i=1}^{w} Z_i = 1]) \geq \sum_{i=1}^{w} log\left(1-\frac{m_i}{2n}\right) \geq -\sum_{i=1}^{w}\frac{m_i}{2n}$$

The fact that $\sum_{i=1}^{w} m_i \leq n$ leads to :

$$\mathbb{P}[\bigwedge_{i=1}^{w} Z_i = 1] \geq \exp\left(-\sum_{i=1}^{w}\frac{m_i}{2n}\right) \geq \exp\left(-\frac{1}{2}\right)$$

The theorem follows. ■

5.2 Probabilistic Gathering

In this section, we analyze the complexity of probabilistic gathering in a fault-free environment. The algorithm presented in this section extends one analyzed in earlier work [1].

We prove that additional information on the environment can significantly improve the convergence time of gathering. Using strong multiplicity knowledge, for example, we obtain a tighter bound of $O(1)$ which significantly improves the best known bound of $O(n^2)$.

In earlier work [1], we proposed a probabilistic algorithm that solves the fault-free gathering in the SYm model, under a specific class of schedulers, known as k-bounded schedulers.[3]

Briefly, the algorithm works as follows. A robot, when chosen by the scheduler, randomly selects one of its neighbors and moves towards its position with probability $\frac{1}{\delta}$, where δ is the size of the robot's view. In the model considered, robots have unlimited visibility, and hence the value of δ is n.

We proved that this strategy probabilistically solves 2-gathering in the *SYm* model under an arbitrary scheduler and converges in 2 rounds in expectation. We also proved that it solves the n-gathering problem ($n \geq 3$), under a fair k-bounded scheduler without multiplicity knowledge and converges under fair bounded schedulers in n^2 rounds in expectation.

[3] In short, a k-bounded scheduler is one ensuring that, during any two consecutive activations of any robot, no other robot is activated more than k times.

We now show that the complexity bound of $O(n^2)$ can be decreased to $O(1)$ when robots use strong multiplicity knowledge. It turns out that the algorithm described above can be modified to meet the latter bound.

Briefly, the modified algorithm (Algorithm 5.2) is based on the following idea. Instead of selecting a destination randomly, a robot p activated by the scheduler moves to the group of robots with maximal multiplicity. When several such groups exist, the robot p tosses a coin and moves only with probability $\frac{1}{2n}$. If it moves, then p arbitrarily selects a point of maximal multiplicity to which it does not belong, and moves toward it.

Algorithm 5.2 Probabilistic gathering executed by robot p with strong multiplicity knowledge.

Functions: $observe\_neighbors$ :: returns the set of robots in the system
$maximal\_multiplicity$ :: returns the set of points with maximal multiplicity
$local\_position$ :: returns the local position of the robot p
(i.e., the locations shared by the largest number of robots as returned by the observe function);

Actions:
$\mathcal{N}_p = observe\_neighbors()$;
if $|maximal\_multiplicity\ (\mathcal{N}_p)| > 1$ **then**
with probability $\frac{1}{2n}$
let q be a point chosen arbitrarily from $maximal\_multiplicity\ (\mathcal{N}_p) \setminus lp()$;
move towards q;
with probability $1 - \frac{1}{2n}$
do not move;
if $|maximal\_multiplicity\ (\mathcal{N}_p)| = 1$ **then**
let q be the point given by $maximal\_multiplicity(\mathcal{N}_p)$;
move towards q;

As we show below, Algorithm 5.2 converges in $O(1)$ rounds in expectation. Interestingly, this means that multiplicity knowledge—used so far in order to break the symmetry of the system—can also considerably help in speeding up the convergence of gathering.

Lemma 1 *Given an arbitrary configuration and only one single robot being activated, the probability to reach a configuration with a unique point of maximal multiplicity is at least* $\frac{1}{2n}$.

Proof: This lemma is an obvious consequence of Algorithm 5.2. Briefly, when there are several points of maximal multiplicity, the activated robot moves with probability $\frac{1}{2n}$ and selects one of the points of maximal multiplicity as its destination. That point's multiplicity then increases by one. ■

We are now able to prove the convergence.

Theorem 2 *Algorithm 5.2 achieves convergence in* $O(1)$ *expected rounds at worst.*

Proof: The analysis can be decomposed into two phases. First, the algorithm must reach a configuration with one single point of maximal multiplicity. Second, the remaining robots must join that point of maximal multiplicity. The second part being obvious, we focus on the first one.

Consider a situation in which, at time t_0, there are at least two different points of maximal multiplicity. We want to obtain an upper bound on the time needed to reach a configuration with only one single point of maximal multiplicity. According to Lemma 1, it is enough to ensure that only one robot changes its position.

Let the period $[t_0, t_0 + k]$ be the shortest fragment where at least n activations occurs. For any $j (0 \leq j \leq k)$, a_j denotes the number of robots activated at time $t_0 + j$. We define X to be the random variable such that $t_0 + X$ is the first time after t_0 when exactly one robot changes its position. We also define Y_j as the indicator random variable of the event that no robot moves during the period $[t_0, t_0 + j - 1]$ (for simplicity of the argument, let $\mathbb{P}[Y_1 = 1] = 1$). Then, we can obtain the following bound

$$\begin{aligned}
\mathbb{P}[X = j \wedge Y_j = 1] &= \binom{a_j}{1}\left(1 - \frac{1}{2n}\right)^{a_j - 1}\left(\frac{1}{2n}\right) \cdot \prod_{h=1}^{j-1}\left(1 - \frac{1}{2n}\right)^{a_h} \\
&= \binom{a_j}{1}\left(1 - \frac{1}{2n}\right)^{a_j - 1}\left(\frac{1}{2n}\right) \cdot \left(1 - \frac{1}{2n}\right)^{\sum_{h=1}^{j-1} a_h} \\
&\geq \left(1 - \frac{a_j - 1}{2n}\right)\left(\frac{a_j}{2n}\right) \cdot \left(1 - \frac{\sum_{h=1}^{j-1} a_h}{2n}\right) \geq \frac{a_j}{8n}
\end{aligned}$$

We use the fact that $a_j \leq n$ and $\sum_{h=1}^{j-1} a_h < n$ (notice that fewer than n robots can be activated by $j-1(<k)$). The above inequality implies the existence of a lower bound for $\mathbb{P}[X=j]$ as follows: $\mathbb{P}[X=j] \geq \mathbb{P}[X=j \wedge Y_j=1] \geq \frac{a_j}{8n}$. Thus, we can bound the probability that the system reaches a configuration with a single point of maximum multiplicity by t_0+k:

$$\mathbb{P}[X \leq k] \quad \geq \quad \sum_{h=1}^{k} \frac{a_j}{8n} \quad \geq \quad \frac{1}{8}$$

Consequently, we can conclude that exactly one point with maximal multiplicity appears with constant probability (more than $\frac{1}{8}$) in a single round. Now, let us consider the two-states variable:

- $Z_t = 1$ when "there is only one maximal multiplicity point at time t."
- $Z_t = 0$ when "there are several maximal multiplicity points at time t."

$(Z_t)_{t\in\mathbb{N}^*}$ is a Markov chain and we need the expectation of the time needed for this stochastic process to reach state 1 starting from state 0. Formally,

$$T_0^1 = \mathbb{E}\left[\min\{t \text{ such that } Z_t = 1 \text{ knowing } Z_0 = 0\}\right]$$

We obtain $T_0^1 = 8$. This is enough to prove that a unique point of maximal multiplicity can be created in $O(1)$ expected rounds. Indeed, this stems from $T_0^1 = 8$ and the fact that as soon as a unique point is created, any activated robot will join it. This also gives that the gathering is achieved in a contant number of rounds. ■

6 The Case of Weak Multiplicity Knowledge

In this section, we assume that robots have *weak multiplicity knowledge*. The algorithms proposed in Sect. 5, based heavily on strong multiplicity knowledge, do not work under weak multiplicity. For the scattering algorithm the multiplicity knowledge is implicitly used in the computation of the number of directions while for the gathering algorithm uses explicitly in the code the strong multiplicity knowledge.

In this section we analyse Algorithm 6.1 (originally presented in [2]) that achieves scattering in a polylogarithmic convergence time with weak multiplicity. For gathering, since it is now impossible for robots to distinguish which point is the point of maximal multiplicity, we instead rely on scattering as an initial step.

6.1 Probabilistic scattering

The algorithm works as follows (Algorithm 6.1). Each robot uses a function $Random()$ that returns a value probabilistically chosen in the set $\{0,1\}$: 0 with probability $\frac{3}{4}$, and 1 with probability $\frac{1}{4}$. When a robot r_i becomes active at time t, it first computes the Voronoi diagram of the set of points occupied by the robots. Then, r_i moves toward an arbitrary location inside its Voronoi cell $Cell_i$ if $Random()$ returns 0. We now look at the convergence time of Algorithm 6.1.

Algorithm 6.1 Probabilistic Scattering executed by robot r_i with weak multiplicity knowledge.

Compute the Voronoi diagram;
$Cell_i :=$ Voronoi cell where r_i is located;
$Current\_Pos :=$ position where r_i is located;

if $Random() = 0$ **then**
    Move toward an arbitrary position in $Cell_i$, different from $Current\_Pos$;
else Do not move;

Lemma 2 *Let a period $[t_0, t_1]$ be the shortest fragment in which all robots are activated at least once and execute Algorithm 6.1. We define R as the set of robots located on a multiplicity point x_0 at time t_0, and Q as the set of points on which at least one robot of R stays at time t_1. Then, with probability at least $1-3e^{-\frac{|R|}{11}}$, all points in Q have multiplicity less than $\frac{3}{4}|R|$.*

Proof: We define the random variable Z as the maximal multiplicity of all points in Q. Therefore, we want to show that $\mathbb{P}[Z \geq \frac{3}{4}|R|] \leq \frac{3}{e^{\frac{|R|}{11}}}$. To do so, we introduce the following two random variables:

- Z_1 is "the multiplicity of point x_0 at time t_1"
- Z_2 is "the maximal multiplicity of all points in $Q \setminus \{x_0\}$ at time t_1"

The reader may easily observe that $\mathbb{P}[Z_1 \geq \frac{3}{4}|R| \wedge Z_2 \geq \frac{3}{4}|R|] = 0$. Thus, we have

$$\mathbb{P}[Z \geq \tfrac{3}{4}|R|] \;=\; \mathbb{P}[Z_1 \geq \tfrac{3}{4}|R| \vee Z_2 \geq \tfrac{3}{4}|R|] \;=\; \mathbb{P}[Z_1 \geq \tfrac{3}{4}|R|] + \mathbb{P}[Z_2 \geq \tfrac{3}{4}|R|]$$

We now study $\mathbb{P}[Z_1 \geq \frac{3}{4}|R|]$ and $\mathbb{P}[Z_2 \geq \frac{3}{4}|R|]$ separately.

Case 1 $\mathbb{P}[Z_1 \geq \frac{3}{4}|R|] \leq e^{-\frac{|R|}{8}}$

Since no point occupied by some robot can become the destination of another robot, no robot moves to x_0 during $[t_0, t_1]$. Thus, if $Z_1 \geq \frac{3}{4}|R|$ holds, at most $\frac{3}{4}|R|$ robots in R keep their position during $[t_0, t_1]$. From the fact that each robot is activated at least once during that period, the probability that some robot on x_0 at t_0 remains on x_0 at t_1 is at most $\frac{1}{2}$. Thus, using the random variable $B_{|R|,\frac{1}{2}}$ following binomial distribution with parameter $|R|$ and $\frac{1}{2}$,[4] we obtain the upper bound for $\mathbb{P}[Z_1 \geq \frac{3}{4}|R|]$:

$$\mathbb{P}[Z_1 \geq \tfrac{3}{4}|R|] \;\leq\; \mathbb{P}[B_{|R|,\frac{1}{2}} \geq \tfrac{3}{4}|R|] \;\leq\; e^{-\frac{|R|}{8}}$$

where we use Chernoff's bound[5] for the binomial distribution $B_{|R|,\frac{1}{2}}$. This proves the case.

Case 2 $\mathbb{P}[Z_2 \geq \frac{3}{4}|R|] \leq 2e^{-\frac{3}{32}|R|}$

To have a point $x \in Q \setminus \{x_0\}$ with multiplicity higher than $\frac{3}{4}|R|$ at t_1, at least $\frac{3}{4}|R|$ robots must leave x_0 by time t_1. In addition, those robots must leave x_0 simultaneously because two robots in R leaving x_0 at different times necessarily have distinct destinations. Thus, there exists a time when more than $\frac{3}{4}|R|$ robots at the same position are activated (if such time does not exist during $[t_0, t_1]$, clearly $\mathbb{P}[Z_2 \geq \frac{3}{4}|R|] = 0$). Let t' be the first time after t_0 when at least $\frac{3}{4}|R|$ robots on x_0 are activated. We define y to be the number of robots on x_0 activated at t', and Y to be the random variable representing the number of robots leaving x_0 at t'. Notice that if the event of $\frac{3}{4}|R| \geq Y \geq \frac{1}{4}|R|$ occurs, the multiplicity of any point in Q becomes less than $\frac{3}{4}|R|$, and thus there is no chance that more than $\frac{3}{4}|R|$ robots located at the same point are simultaneously activated after t', that is, $\mathbb{P}[\{Z_2 \geq \frac{3}{4}|R|\} \mid \{\frac{3}{4}|R| \geq Y \geq \frac{1}{4}|R|\}] = 0$. Thus, it results that $\mathbb{P}[Z_2 \geq \frac{3}{4}|R|] \leq \mathbb{P}[\{\frac{3}{4}|R| \leq Y\} \vee \{\frac{1}{4}|R| \geq Y\}]$ holds. Since Y is equivalent to the binomial random variable $B_{y,\frac{1}{2}}$, using the same method as in the previous case, we obtain the following bound, which proves the case:

$$\begin{aligned}\mathbb{P}[Z_2 \geq \tfrac{3}{4}|R|] &\leq \mathbb{P}[\left\{\frac{3}{4}|R| \leq Y\right\} \vee \left\{\frac{1}{4}|R| \geq Y\right\}] \\ &\leq (\mathbb{P}[B_{y,\frac{1}{2}} \geq \frac{3}{4}|R|] + \mathbb{P}[B_{y,\frac{1}{2}} \leq \frac{1}{4}|R|]) \;\leq\; 2e^{-\frac{y}{8}} \;\leq\; 2e^{-\frac{3}{32}|R|}\end{aligned}$$

Consequently, we obtain $\mathbb{P}[Z \geq \frac{3}{4}|R|] \leq e^{-\frac{|R|}{8}} + 2e^{-\frac{3}{32}|R|} \leq 3e^{-\frac{|R|}{11}}$. ■

Lemma 3 *Let a time interval $[t_0, t_1]$ be a period containing k rounds, and let m be the maximal multiplicity at time t_0. Then, the maximal multiplicity at time t_1 is less than $\frac{3m}{4}$ with probability $(1 - \frac{4n}{me^{\frac{km}{11}}})$.*

Proof: Let $\mathcal{P} = \{p_1, p_2, \cdots, p_w\}$ be the set of points occupied by more than $\frac{3m}{4}$ robots at t_0, where $w = |P|$. We define the indicator random variable Z_i such that $Z_i = 1$ if p_i is divided into two or more points with multiplicity less than $\frac{3m}{4}$, and $Z_i = 0$ otherwise. Again, thanks to Voronoi diagrams, $\{Z_1, Z_2, \cdots\}$ are mutually independent.

[4] That is, the probability distribution of $B_{l,q}$ is defined as $\mathbb{P}[B_{l,q} = k] = \binom{k}{l} q^k (1-q)^{l-k}$ for any $k(0 \leq k \leq l)$.

[5] For any $\delta > 0$, $\mathbb{P}[B_{l,q} \geq (1+\delta)lq] \leq e^{-\delta^2 lq}$.

From Lemma 2, the probability that the respective multiplicity of each point in W decreases to less than $\frac{3m}{4}$ is bounded:

$$\begin{aligned}\mathbb{P}\left[\bigwedge_{i=1}^{w}(Z_i=1)\right] &= \prod_{i=1}^{w}\mathbb{P}[Z_i=1] \geq (1-3e^{-\frac{km}{11}})^w \\ &\geq (1-3we^{-\frac{km}{11}}) \geq \left(1-\frac{4n}{me^{\frac{km}{11}}}\right)\end{aligned}$$

where we use the fact that $w \cdot \frac{3m}{4} \leq n$. This completes the proof. ■

We deduce from Lemma 3 the following two corollaries.

Corollary 1 *The maximal multiplicity decreases to less than* $11 \log 4n$ *within* $O(\log n)$ *rounds in expectation.*

Corollary 2 *Starting from the configuration with maximal multiplicity* $11 \log 4n$ *or less, the scattering is achieved within* $O(\log n \log(\log n))$ *rounds in expectation.*

The first corollary follows from Lemma 3 with $k = 1$ for the case of $m > 11 \log 4n$, and the second one follows from the lemma with $k = 11 \log 4n$ for $m \leq 11 \log 4n$. By combining both corollaries, we obtain the convergence of Algorithm 6.1.

Theorem 3 *The convergence time of Algorithm 6.1 is* $O(\log n \log(\log n))$.

6.2 Probabilistic gathering

We now present an algorithm for probabilistic gathering under weak multiplicity knowledge. Algorithm 6.2, is built similarly to the one presented in Section 5.2 for strong multiplicity, except that it now relies on scattering (Algorithm 6.1) as an initial step.

Algorithm 6.2 Probabilistic gathering executed by robot p with weak multiplicity knowledge.

```
Actions:   N_p = observe_neighbors();
           if N_p includes two or more multiplicity points then
               Execute Algorithm 6.1;
           else if N_p includes a single multiple point q then
               Move towards q;
           else
               let q be a point chosen arbitrarily from N_p;
               With probability 1/(2n)
                   move towards q;
               With probability 1 - 1/(2n)
                   do not move;
```

Lemma 4 *Starting from any scattered configuration, a configuration with exactly one multiplicity point is reached in one round with constant probability.*

The proof is identical to that of Theorem 2.

Since scattering is achieved within $O(\log n \log(\log n))$ rounds in expectation, the above lemma implies that the system can have exactly one multiplicity point within $O(\log n \log(\log n))$ rounds, starting from an arbitrary configuration. Thus, the convergence time of Algorithm 6.2 directly follows.

Theorem 4 *The convergence time of Algorithm 6.2 is* $O(\log n \log(\log n))$.

7 Conclusion and Discussions

The contribution of this paper is twofold. First, we proposed a detailed complexity analysis of existing probabilistic agreement algorithms (gathering and scattering). Second, using Markov chains tools and multiplicity knowledge, we proved that the convergence time of gathering can be reduced from $O(n^2)$ (the best-known bound so far) to $O(1)$ or $O(\log n \cdot \log(\log n))$, depending on the model of multiplicity. We have also shown the same bounds for the scattering problem, thus closing the complexity gap between the two problems (gathering and scattering are two important agreement problems in robots networks).

Our work also confirms that the analysis based on Markov chains fits well with the oblivious robot networks model. Therefore, we intend to apply this approach to analyze other algorithms for oblivious robot networks, such as, leader election or flocking.

Acknowledgments

We would like to express our gratitude to the associate editor and the anonymous reviewers for their numerous insightful comments that greatly helped improve this manuscript. This work is partially supported by Région Île-de-France (Julien Clement), by MEXT Grant-in-Aid for Young Scientists (A) No. 18680007 (Xavier Défago), and by the Telecommunication Advancement Foundation (Taisuke Izumi).

References

[1] Xavier Défago, Maria Gradinariu, Stéphane Messika, and Philippe Raipin Parvédy. Fault-tolerant and self-stabilizing mobile robots gathering. In Shlomi Dolev, editor, *DISC*, pages 46–60. Springer, 2006.

[2] Yoann Dieudonné and Franck Petit. Robots and demons (code of the origins). In *FUN*, pages 108–119, 2007.

[3] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, San Francisco, CA, USA, 1996.

[4] N. A. Lynch, R. Segala, and F. W. Vaandrager. Hybrid I/O automata. *Information and Computation*, 185(1):105–157, August 2003.

[5] J.R. Norris. *Markov Chains*. Cambridge University Press, 1997.

[6] G. Prencipe. CORDA: Distributed coordination of a set of autonomous mobile robots. In *Proc. 4th European Research Seminar on Advances in Distributed Systems (ERSADS'01)*, pages 185–190, May 2001.

[7] I. Suzuki and M. Yamashita. A theory of distributed anonymous mobile robots formation and agreement problems. Technical report, Wisconsin Univ. Milwakee, Dept. of EECS, 1994.

[8] I. Suzuki and M. Yamashita. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM Journal of Computing*, 28(4):1347–1363, 1999.