



Continuous Protocols for Swarm Robotics

Peter Kling^{1(✉)} and Friedhelm Meyer auf der Heide²

¹ Universität Hamburg, Hamburg, Germany
peter.kling@uni-hamburg.de

² University of Paderborn, Paderborn, Germany
fmadh@uni-paderborn.de

Abstract. We consider simple models of swarms of identical, anonymous robots: they are points in the plane and “see” only their neighbors (robots within distance one). We will deal with distributed local protocols of such swarms that result in formations like “gathering at one point”. The focus will be on protocols assuming a continuous time model. We present upper and lower bounds on their run time and energy consumption, and compare different protocols both theoretically and experimentally.

Keywords: Robots · Continuous · Gathering

1 Introduction

Envision a scenario where n mobile robots, each having a limited viewing range, are placed in the Euclidean plane and are supposed to establish a certain formation. To reach this formation, each robot has to plan and perform its movement based solely on the positions of the other robots within its viewing range, which we normalize to 1. In particular, the robots are not provided with global view, communication, or long term memory.

This chapter considers one of the most basic formation problems: the GATHERING problem. Here, the n robots must move such that, eventually, they gather at a single, not predetermined point. While performing their protocols, it is crucial that the visibility graph spanned by the robots stays connected; if a robot loses sight of all other robots, no deterministic, local protocol can be guaranteed to reconnect the lost robot to the remaining formation.

Most protocols for such formation problems are based on some kind of discrete round model. For example, Ando et al. [1] and Degener et al. [4] show that gathering can be achieved with a simple protocol by robots with a limited visibility in a synchronous, discrete time model in $O(n^2)$ rounds. For the same problem with an unlimited viewing range, Cohen and Peleg [3] analyze a simple algorithm in several asynchronous time models. Further publications consider the GATHERING problem in similar time models, see for example [2, 6, 9–11, 16, 17].

All of the above models have in common that they are based on the so-called Look-Compute-Move (*LCM*) model. That is, the robots act in rounds, where each

round consists of a **Look** operation, a **Compute** operation, and a **Move** operation. During the **Look** operation, a robot determines the positions of all visible robots in its vicinity. During the **Compute** operation, the observed information is used to determine a target point. Finally, during the **Move** operation, the robot moves towards the previously computed target point. The specific models differ in whether these operations are executed synchronously or asynchronously (or something in between).

A different approach is to use a continuous time model. This was first done by Gordon et al. [8] for the GATHERING problem. In such a continuous time model, all robots perpetually and at the same time measure and adjust their movement paths. This causes the trajectories of the robots, who are assumed to have some constant maximum moving speed, to become (continuous) curves. While this continuous motion model is somewhat idealized and might seem unrealistic – given that we assume there is no delay between the robots’ sensors and actors – it is comparatively close to real applications [14].

This chapter presents some of the more recent results that introduced general techniques to analyze the time required by such distributed robot formation protocols in the continuous time model to reach their goal and discusses other aspects specific to this model.

Chapter Outline. We continue with a formal model description and with the introduction of continuous robot formation protocols in Sect. 2. Afterward, we collect some auxiliary results in Sect. 3 that will be used throughout the chapter. In Sect. 4 we introduce a general class of (not necessarily distributed) robot formation protocols, so-called *contracting protocols*, and study their gathering time. Section 5 introduces a specific protocol of this class and sketches the proof that this protocol solves the gathering problem in asymptotical optimal time. We conclude this chapter in Sect. 6, where we discuss the issue of collisions and how to avoid them.

2 Model Description and Continuous Protocols

Consider a set of n autonomous, mobile robots in the Euclidean plane \mathbb{R}^2 . The robots have no spatial dimension and each of them has its own, local coordinate system. In particular, there is no common “origin” notion and no common directional notions like “left” or “right”. All robots have a *visual range* of 1, allowing them to perceive other robots that are within this distance and to determine the relative position of such robots. However, neither can robots distinguish other robots from one another nor is there any form of multiplicity detection: a robot can only distinguish whether there is either *no* robot or *at least one* robot at a given position. There is a continuous notion of time and we assume robots move with a *maximum speed* of 1. We assume an idealized sensor-actor mechanism, allowing a robot to act instantly on any data perceived from within its visual range.

In order to convey the basics and characteristics of continuous robot formation problems (in contrast to their more standard discrete time pendants), we concentrate on the probably most basic formation problem: *gathering*.

Definition 1 (GATHERING Problem). *Consider n autonomous, mobile robots in the Euclidean plane. In the GATHERING problem, we seek to gather all robots at a single point.*

A (*gathering*) protocol is an algorithmic description that takes the current robot positions and determines how each robot moves. We are mostly interested in *distributed protocols*, where each robot determines its movement by the same algorithm that uses only information available to the respective robot (basically the relative positions of other robots within its visual range). The quality of a protocol is measured by its *gathering time*, the worst-case time (over the set of all possible initial robot positions) required by the protocol to gather all robots.

General Notation. Before we dive deeper into the specifics of continuous robot formation protocols, we introduce some general notation. For an integer $m \in \mathbb{N}$ define $[m] := \{1, 2, \dots, m\}$. For $x \in \mathbb{R}^2$ we use $\|x\|_2$ to denote the Euclidean norm (vector length) of x . For $S_1, S_2 \subseteq \mathbb{R}$ and $c_1, c_2 \in \mathbb{R}$ define $c_1 \cdot S_1 + c_2 \cdot S_2 := \{c_1 \cdot s_1 + c_2 \cdot s_2 \mid s_1 \in S_1, s_2 \in S_2\}$. For $x, y \in \mathbb{R}^2$ let $\angle(x, y) \in [-\pi, \pi]$ denote the signed¹ angle formed by the vectors x and y .

Fix an arbitrary ordering of the n robots. A *configuration* $c = (c_i)_{i=1}^n$ is a vector whose i -th element $c_i \in \mathbb{R}^2$ specifies the positions of the i -th robot. Define the polygon $\text{CH}_c \subseteq \mathbb{R}^2$ as the convex hull of all robot positions in configuration c . Furthermore, let $\text{Bord}_c := \{i \mid c_i \in \partial\text{CH}_c\}$ be the set of robots that are at the boundary of the configuration’s convex hull. Similarly, let $\text{Corn}_c \subseteq \text{Bord}_c$ be the set of robots that are at a vertex of the polygon CH_c . A robot from Bord_c is called *border robot* and a robot from Corn_c is called *corner robot*. The *diameter* $\Delta_c := \max \{\|c_i - c_j\|_2 \mid i, j \in [n]\}$ of configuration c is the maximum distance between any two robots.

The *visibility graph* $G_c = (V_c, E_c)$ of configuration c is the Euclidean graph whose vertices are the robots’ positions and in which two vertices are connected by an edge if and only if the two corresponding robots are within viewing range of each other. More formally, $V_c := \{c_i \mid i \in [n]\}$ and $E_c := \{\{u, v\} \mid u, v \in V_c, u \neq v, \|u - v\|_2 \leq 1\}$. A configuration is called *connected* if G_c is connected. We use \mathcal{C}_n to denote the set of all connected configurations of n robots.

Continuous Protocols. A continuous robot formation protocol \mathcal{P} specifies how, at any time $t \geq 0$, each robot calculates its *velocity vector*, which dictates the robot’s current speed and its movement direction. Given a continuous robot formation protocol \mathcal{P} , let $c(t) = (c_i(t))_{i=1}^n$ be the configuration at time t , such that the function $c_i: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^2$ is the trajectory of the i -th robot. For $i \in [n]$ and $t \geq 0$ let $v_i(t)$ denote the velocity vector of the i -th robot at time t .

¹ Without loss of generality, we say an angle is positive if it is measured counterclockwise and negative if it is measured clockwise.

That is, at time t robot i moves with speed $\|v_i(t)\|_2$ in direction $v_i(t)$. We use $s_i(t) := \|v_i(t)\|_2$ as a shorthand for the speed of robot i at time t .

The trajectories c_i are continuous but not necessarily differentiable. Indeed, robots can change their speed and movement direction non-continuously, resulting in a non-differentiable trajectory. However, natural protocols have right-differentiable trajectories, and we restrict our study to such protocols. In particular, this allows us to see robot i 's velocity vector $v_i: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^2$ as the (right) derivative of c_i and we can write $v_i = \dot{c}_i$, where the differentiation is understood to be a right derivative whenever necessary.

It will be useful to consider how robots move relative to each other. For this purpose, we define the angles $\beta_{i,j}(t) := \angle(v_i(t), c_j(t) - c_i(t))$. That is, $\beta_{i,j}(t)$ is the signed angle between the velocity vector of robot i and the line segment connecting robot i and robot j .

Figure 1 illustrates the notions introduced above.




Fig. 1. (a) An example configuration c for $n = 13$ robots with the corresponding visibility graph G_c and the convex hull CH_c . The convex hull has 7 border robots, 6 of which are also corner robots. (b) The current movement vector $v_i(t)$ of a robot i , its current relative positional angle $\beta_{i,j}(t)$ with respect to robot j and example trajectories (in blue). (Color figure online)

3 Auxiliary Results

This section collects some basic results that turn out to be useful in the analysis of continuous robot formation protocols. We start with some simple trigonometric inequalities.

Lemma 2

- (a) For $\alpha \in [0, \pi/2]$ we have $\cos \alpha \geq 1 - 2\alpha/\pi$.
- (b) For $\alpha \geq 0$ we have $\cos \alpha \geq 1 - \alpha^2/2$.
- (c) For $\phi \in [0, 1]$ and $\alpha \in [0, \pi]$ we have $\cos(\phi \cdot \alpha) + \cos((1-\phi) \cdot \alpha) \geq 2 \cdot (1 - \alpha/\pi)^2$.

Proof

- (a) The statement follows from basic calculus by realizing that the graph of $f(x) = 1 - 2x/\pi$ is a secant line that intersects the graph of $\cos(x)$ at $x = 0$ and $x = \pi/2$.
- (b) The statement follows by realizing that $\cos \alpha$ and $1 - \alpha^2/2$ are equal for $\alpha = 0$ and that the first term's derivative ($-\sin \alpha$) is as least as large as the second term's derivative ($-\alpha$) for any $\alpha \geq 0$.
- (c) For $x, y \in \mathbb{R}$ we have the trigonometric identity $\cos x + \cos y = 2 \cdot \cos((x+y)/2) \cdot \cos((x-y)/2)$ (see [15, Identity 4.21.8]). We apply this to the left-hand side of the desired inequality and calculate

$$\begin{aligned} \cos(\phi \cdot \alpha) + \cos((1 - \phi) \cdot \alpha) &= 2 \cdot \cos\left(\frac{\alpha}{2}\right) \cdot \cos\left(\frac{2\phi - 1}{2} \cdot \alpha\right) \\ &\geq 2 \cdot \left(\cos\left(\frac{\alpha}{2}\right)\right)^2 \geq 2 \cdot \left(1 - \frac{\alpha}{\pi}\right)^2. \end{aligned} \quad (1)$$

The first inequality uses $(2\phi - 1)/2 \cdot \alpha \in [-\alpha/2, \alpha/2] \subseteq [-\pi/2, \pi/2]$. Thus, the cosine in the expression is minimized when $(2\phi - 1)/2 \cdot \alpha = \alpha/2$. The second inequality uses the lemma's first statement. \square

The next lemma is central for analyzing the gathering time of protocols. Our general progress measure is based on how the robots' convex hull changes over time. We quantify this change by studying how the distance between neighboring corner robots changes. To this end, the following lemma expresses the change in the distance between two robots i and j in terms of their current speeds $s_i(t)$ and $s_j(t)$ and their direction of movement relative to each other (the angles $\beta_{i,j}(t)$ and $\beta_{j,i}(t)$ between their velocity vectors and their connecting line). See Fig. 2 for an illustration.

Lemma 3. *Consider a robot formation protocol \mathcal{P} and fix two robots i and j . The distance $d(t) := \|c_i(t) - c_j(t)\|_2$ between i and j at time t changes with speed*

$$\dot{d}(t) = -s_i(t) \cdot \cos \beta_{i,j}(t) - s_j(t) \cdot \cos \beta_{j,i}(t). \quad (2)$$

Proof. Define $D: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^2, t \mapsto c_j(t) - c_i(t)$, such that $d(t) = \|D(t)\|_2$. We use $D_x(t)$ to refer to the x -component of $D(t)$ and, similarly, $D_y(t)$ to refer to the y -component of $D(t)$. Fix a time $t \geq 0$. Without loss of generality, we can translate and rotate the coordinate system such that $D(t) = (d(t), 0)$. This immediately yields

$$\dot{d}(t) = \left(\frac{D_x(t)}{d(t)}, \frac{D_y(t)}{d(t)} \right) \cdot \begin{pmatrix} \dot{D}_x(t) \\ \dot{D}_y(t) \end{pmatrix} = \dot{D}_x(t). \quad (3)$$

Note that $\dot{D}_x(t)$ is the x -component of $\dot{D}(t)$, which can be written as

$$\begin{aligned} \dot{D}(t) &= \dot{c}_j(t) - \dot{c}_i(t) = v_j(t) - v_i(t) \\ &= s_j(t) \cdot (-\cos \beta_{j,i}(t), \sin \beta_{j,i}(t)) - s_i(t) \cdot (\cos \beta_{i,j}(t), \sin \beta_{i,j}(t)). \end{aligned} \quad (4)$$

Together, Eqs. (3) and (4) imply the lemma's statement. \square




Fig. 2. Illustration of Lemma 3

4 Contracting Robot Formation Protocols

A natural property of protocols that solve the GATHERING problem is that robots move “towards each other”, causing the convex hull of all robot positions to contract over time. In the following we define the class of *contracting* robot formation protocols – which formalizes this intuitive property – and prove general upper and lower bounds for this class.

Define $p(t) := |\{c_i(t) \mid i \in \text{Corn}_{c(t)}\}|$ as the number of vertices of the polygon $\text{CH}_{c(t)}$. Let $m_1(t), m_2(t), \dots, m_{p(t)}(t)$ denote the vertices ordered counter-clockwise along the boundary of $\text{CH}_{c(t)}$ (starting at an arbitrary vertex). For convenience, define $m_0(t) := m_{p(t)}(t)$ and $m_{p(t)+1}(t) := m_1(t)$.

Definition 4 (Length). *The length $l(t)$ of the configuration at time t is*

$$l(t) := \sum_{i=1}^{p(t)} \|m_i(t) - m_{i-1}(t)\|_2. \tag{5}$$

We use the shorthand $l := l(0)$ to denote the length of the initial configuration.

Note that the robots have solved the GATHERING problem at time t if and only if $l(t) = 0$. This property (and the fact that most reasonable protocols do not increase a configuration’s length) make the length a good way to measure the progress of a gathering protocol. Next we define a quite general class of (not necessarily distributed) robot formation protocols. The definition, which is originally from [13], simply requires that corner robots move with maximum speed in some direction within the robots’ convex hull.

Definition 5 (Contracting [13]). *Consider a continuous robot formation protocol \mathcal{P} . We say \mathcal{P} is contracting if for any time $t \geq 0$ with $l(t) > 0$ each corner robot $i \in \text{Corn}_{c(t)}$ moves with speed $s_i(t) = 1$ along the velocity vector $v_i(t) \in (m_{i-1}(t) - m_i(t)) \cdot \mathbb{R}_{\geq 0} + (m_{i+1}(t) - m_i(t)) \cdot \mathbb{R}_{\geq 0}$.*

See Fig. 3 for an illustration.




Fig. 3. Possible movement vectors for a contracting protocol. Note that the velocity vectors of corner robots have all the same length (1), while the border robot may move at a slower speed. The movement of robots within the convex hull is completely unrestricted.

Even though Definition 5 does not specify the movement of general border robots, it has an interesting implication for their movement. Consider a border robot $j \in \text{Bord}_{c(t)} \setminus \text{Corn}_{c(t)}$ between two corner robots $i_1, i_2 \in \text{Corn}_{c(t)}$ at time t of a contracting robot formation protocol \mathcal{P} . Assume that j “falls behind” at time t , such that for any small enough $\epsilon > 0$ j is a corner robot (at a position different from i_1 and i_2) in configuration $c(t + \epsilon)$. Definition 5 requires that j moves with speed 1 towards the line connecting i_1 and i_2 for any such ϵ . In the worst case this line moves with a speed of at most 1 away from j (if both i_1 and i_2 move accordingly). Thus, the distance between j and the line connecting i_1 and i_2 cannot increase at time $t + \epsilon$. But if that holds for any $\epsilon > 0$, j cannot leave the line in the first place. In other words, contracting protocols will not create new corners along the robots’ convex hull.

The above phenomenon illustrates an important aspect of continuous strategies and is also known as *Zenoness* (see [8]). The continuous nature of the system – which allows for an instant and continuous course correction – makes it possible that collinear robots (like corner robots and any border robots between them) remain collinear.

4.1 Gathering Time of Contracting Protocols

With the above notions we can formulate a general result that bounds the gathering time of *any* contracting robot formation protocol.

Theorem 6 ([13]). *Consider a continuous robot formation protocol \mathcal{P} started in a configuration of n robots and of diameter Δ . If \mathcal{P} is contracting, then it has gathering time at most $\pi \cdot n \cdot \Delta/8$.*

Proof. Fix a time t and consider the convex hull $\text{CH}_{c(t)}$ formed by the robots’ positions at time t . The robots have gathered at time t if and only if the length $l(t)$ of the robots’ convex hull equals zero. We show that at any time t with

$l(t) > 0$, the configuration's length decreases with a speed of at least $8/n$. Since the initial length is at most $2\pi \cdot \Delta/2$, this implies the theorem's statement.

So fix a time $t \geq 0$ with $l(t) > 0$. Recall the definition of the vertices $m_1(t), m_2(t), \dots, m_{p(t)}(t)$ of the robots' convex hull and of the configuration's length $l(t) = \sum_{\iota=1}^{p(t)} \|m_\iota(t) - m_{\iota-1}(t)\|_2$. To avoid double indices, we make a slight abuse of notation and identify ι with one of the robots positioned on $m_\iota(t)$. In particular, we write $v_\iota(t)$ for the velocity vector of the robots positioned at $m_\iota(t)$. Similarly we write $\beta_{\iota,\iota-1}(t)$ and $\beta_{\iota-1,\iota}(t)$ for the corresponding angles between robots positioned at $m_\iota(t)$ and $m_{\iota-1}(t)$. Note that, since the robot formation protocol \mathcal{P} is contracting, these vertices move with speed $s_\iota(t) = 1$. Define $d_\iota(t) := \|m_\iota(t) - m_{\iota-1}(t)\|_2$ as the distance between the corner robots at $m_\iota(t)$ and $m_{\iota-1}(t)$. By Lemma 3

$$\dot{d}_\iota(t) = -\cos \beta_{\iota,\iota-1}(t) - \cos \beta_{\iota-1,\iota}(t). \tag{6}$$

For $\iota \in [p(t)]$ let $\alpha_\iota(t) \in [0, \pi]$ denote the inner angle of the polygon $\text{CH}_{c(t)}$ at vertex $m_\iota(t)$. Since the robot formation protocol \mathcal{P} is contracting, the velocity vector $v_\iota(t)$ points towards the inside of the robots' convex hull, such that $\alpha_\iota(t) = \beta_{\iota,\iota-1}(t) + \beta_{\iota,\iota+1}(t)$. Since $l(t) = \sum_{\iota=1}^{p(t)} d_\iota(t)$, we can take the derivative of $l(t)$ and apply Eq. (6) to get

$$\begin{aligned} \dot{l}(t) &= \sum_{\iota=1}^{p(t)} (-\cos \beta_{\iota,\iota-1}(t) - \cos \beta_{\iota-1,\iota}(t)) \\ &= -\sum_{\iota=1}^{p(t)} (\cos \beta_{\iota,\iota-1}(t) + \cos \beta_{\iota,\iota+1}(t)). \end{aligned} \tag{7}$$

Applying Lemma 2(c) to the last expression yields

$$\dot{l}(t) \leq -2 \cdot \sum_{\iota=1}^{p(t)} \left(1 - \frac{\alpha_\iota(t)}{\pi}\right)^2 = -\frac{2}{\pi^2} \cdot \sum_{\iota=1}^{p(t)} (\pi - \alpha_\iota(t))^2. \tag{8}$$

Now we first use the Cauchy-Schwarz inequality and then that the sum of internal angles of a polygon with p vertices equals $(p - 2) \cdot \pi$ to get

$$\begin{aligned} \dot{l}(t) &\leq -\frac{2}{p(t) \cdot \pi^2} \cdot \left(\sum_{\iota=1}^{p(t)} (\pi - \alpha_\iota(t))\right)^2 \\ &= -\frac{2}{p(t) \cdot \pi^2} \cdot (p(t) \cdot \pi - (p(t) - 2) \cdot \pi)^2 \\ &\leq -\frac{8}{p(t)} \leq -\frac{8}{n}. \end{aligned} \tag{9}$$

As argued at the beginning of the proof, this yields the desired statement. \square

In Theorem 6 we did not restrict the configuration to be connected. As long as the protocol (which could be executed by an omniscient observer that knows where to move the robots) is contracting, the stated bound holds. If we start in a connected configuration, as required for any truly distributed protocol restricted by the robots' visual range, we see that the initial diameter is at most $n-1$, yielding a bound of $\pi/8 \cdot n^2$. In fact, if instead of bounding the initial configuration's length by $2\pi \cdot \Delta/2$ in the proof of Theorem 6 we use the bound $l(0) \leq 2(n-1)$ (which can easily be shown, see [12]), we get the following corollary:

Corollary 7 ([13]). *Consider a continuous robot formation protocol \mathcal{P} started in a connected configuration of n robots. If \mathcal{P} is contracting, then it has gathering time at most $n^2/4$.*

4.2 Worst-Case Contracting Protocols

It is not difficult to see, that the upper bound of $O(n^2)$ from Corollary 7 is tight in the sense that there are contracting protocols and corresponding initial configurations for which the gathering time is *at least* $\Omega(n^2)$. In fact, if we assume that all robots are positioned at corners of a regular polygon and move towards their counterclockwise neighbor, we get a situation that resembles the so-called *n-bugs problem* [18], which is known to have a quadratic convergence speed. The next lemma provides this result and a simple proof using our terminology.


Lemma 8. *There is a contracting robot formation protocol \mathcal{P} and a connected initial configuration $c \in \mathcal{C}_n$ such that the gathering time is at least $n^2/(2\pi^2)$.*

Proof. Assume the initial configuration is such that the robots' convex hull is an n -sided regular polygon with edge length 1 (the robots' visual range) and assume that each robot moves with speed 1 towards its counterclockwise neighbor. Since all robots move symmetrically, the configuration will stay an n -sided regular polygon (which eventually degenerates to a point) at any time $t \geq 0$. Now fix a time $t \geq 0$ with $l(t) > 0$. Consider Eq. (7) from the proof of Theorem 6, which gives the speed at which the configuration's length changes. Since the current convex hull is an n -sided regular polygon, we have for each $\iota \in \{1, 2, \dots, p(t)\}$ that $\beta_{\iota, \iota+1} = 0$ and $\beta_{\iota, \iota-1} = (n-2) \cdot \pi/n = \pi - 2\pi/n$. With this, Eq. (7) simplifies to

$$\begin{aligned} \dot{l}(t) &= \sum_{\iota=1}^n -(\cos(\pi - 2\pi/n) + \cos 0) \\ &= -n \cdot (1 - \cos(2\pi/n)) \\ &\geq -n \cdot 2\pi^2/n^2 = -2\pi^2/n, \end{aligned} \tag{10}$$

where the inequality uses Lemma 2(b). Since the initial configuration has length n , this yields the lemma's claim. \square

See Fig. 4 for an illustration of Lemma 8. When introducing the MOVE-ON-BISECTOR protocol in the next section, we will also encounter a "best-case" contracting protocol (global MOVE-ON-BISECTOR).



(a) The robots' velocity vectors for $n = 7$. (b) The robots' trajectories for $n = 10$.

Fig. 4. Configuration, velocity vectors, and trajectories of the n -bug problem from Lemma 8.

5 Near-Optimal Continuous Protocols for Gathering


Section 4 showed that any contracting robot formation protocol has at most quadratic gathering time and that this bound is tight in the sense that there are contracting protocols (and corresponding initial configurations) that have at least quadratic gathering time. However, there is still hope that a specific contracting protocol can have a much better – maybe even linear – gathering time. This is indeed true and has been proved for the so-called MOVE-ON-BISECTOR protocol in [5]. This section provides a sketch of this result; see [5] for the full proof (and further related results).

5.1 The MOVE-ON-BISECTOR Protocol

At each time t , each robot i observes all positions of its neighbors (i.e., other robots within its visual range). It then computes the *local convex hull* $\text{CH}_i(t)$ of its own and any observed positions. Given this local convex hull, robot i performs the following actions:

- (a) If i is on a vertex of $\text{CH}_i(t)$, it moves with speed 1 along the angle bisector of this vertex.
- (b) If i is not on a vertex but on the boundary $\partial\text{CH}_i(t)$ of the local convex hull, it moves with the corresponding line such that it stays on it and maintains the ratio of distances between its two neighbors on the boundary.
- (c) If i is strictly inside $\text{CH}_i(t)$, it does not move.

This protocol was originally suggested by Gordon et al. [8] (using a slightly different description and not under this name). See Fig. 5a for an illustration.



(a) MOVE-ON-BISECTOR protocol with the local convex hull of a moving (blue) and a not moving (green) node.

(b) Global MOVE-ON-BISECTOR protocol for the same configuration.

Fig. 5. (Local) MOVE-ON-BISECTOR vs. global MOVE-ON-BISECTOR. Note that the lengths of the velocity vectors is not accurate. (Color figure online)

Global MOVE-ON-BISECTOR – A Best-Case Contracting Protocol.

Before we analyze the MOVE-ON-BISECTOR protocol, let us give an intuition why it should be fast. Remember that our definition of contracting robot formation protocols is not restricted to *distributed* protocols. So, consider a global variant of MOVE-ON-BISECTOR, where the only difference in the protocol description is that robots use the global convex hull $\text{CH}_{c(t)}$ instead of their respective local convex hull $\text{CH}_i(t)$. This is obviously not a distributed protocol, as robots do not have, in general, knowledge of the global convex hull. However, it is a contracting protocol (see also Fig. 5b). In fact, in a sense this is an “optimal” contracting gathering protocol: Using an analysis similar to Sect. 4.1, one can see that, as long as the robots have not yet gathered, the length of the configuration at time t decreases at a constant rate, yielding a gathering time that is linear in the diameter Δ of the initial configuration. Since robots move at unit speed, this is asymptotically optimal.

Unfortunately, we cannot analyze the actual MOVE-ON-BISECTOR protocol in the same way as its global counterpart. Since robots move not along the vertices of the global convex hull, the length of the configuration does not necessarily decrease at a constant speed. However, the next section introduces a “local” variant of a configuration’s length, which we call *stretch*. With some additional work, we can show (Sect. 5.3) that not only does this stretch decrease at a constant rate but it is also linear in the configuration’s diameter. This yields the same asymptotical optimal guarantee as the global MOVE-ON-BISECTOR protocol.

5.2 Preliminaries

Instead of using the robots’ convex hull and its circumference as a progress measure, our progress measure for the MOVE-ON-BISECTOR protocol is based

on the simple polygon spanned by the robots (the *configuration polygon*) and its circumference (its *stretch*).

Definition 9 (Configuration Polygon). *Consider a configuration c with its visibility graph G_c . This graph defines a simple polygon Pol_c . We call this polygon the configuration polygon of configuration c .*

Note that the configuration polygon or parts of it might be degenerated (e.g., to a line), as can be seen in the example from Fig. 6.

Define $q(t)$ as the number of vertices of the configuration polygon $\text{Pol}_{c(t)}$. Let $w_1(t), w_2(t), \dots, w_{q(t)}(t)$ denote the vertices ordered counterclockwise along the boundary of $\text{Pol}_{c(t)}$ (starting at an arbitrary vertex). For convenience, define $w_0(t) := w_{q(t)}(t)$ and $w_{q(t)+1}(t) := w_1(t)$. Similar to the length of the configuration we can define the *stretch* of a configuration as follows:

Definition 10 (Stretch). *The stretch $\mathfrak{s}(t)$ of the configuration at time t is*

$$\mathfrak{s}(t) := \sum_{\iota=1}^{q(t)} \|w_\iota(t) - w_{\iota-1}(t)\|_2 \tag{11}$$

We use the shorthand $\mathfrak{s} := \mathfrak{s}(t)$ to denote the stretch of the initial configuration.

For a vertex $w_\iota(t)$ of the configuration polygon at time t we define $\alpha_\iota(t) \in [0, 2\pi)$ as the *inner* angle of the configuration polygon at that vertex. We use $\mathcal{W}(t) := \{\iota \in [q(t)] \mid \alpha_\iota(t) < \pi\}$ to characterize the set of all convex angles of the polygon and $\bar{\mathcal{W}}(t) := [q(t)] \setminus \mathcal{W}(t)$ for the remaining (concave) angles. See Fig. 6 for an illustration of the notions defined above.

An important observation is, that a robot at the vertex w_ι of the configuration polygon at time t is not necessarily on the vertex of its local convex hull. Indeed, if the inner angle at that vertex is concave ($> \pi$), it may be inside its local convex hull and may, thus, not move. However, any robot at a vertex that forms a convex inner angle is guaranteed to move, as they must be on a vertex of their local convex hull. These are exactly the vertices w_ι with $\iota \in \mathcal{W}$.

Note that while the trajectory of the robots is continuous, the direction in which a robot moves may change in a non-continuous way. This may happen if the visibility graph changes. Moreover, a change in the visibility graph may also influence the length of the current configuration in a non-continuous way. However, it is not hard to see (and it has been proved in [8]) that once two robots are within visual range, they will not lose visibility. Thus, there can be only a finite amount of these discontinuities. Moreover, by the definition of the visibility graph, the number of the configuration polygon’s vertices cannot increase at such a discontinuity. This leads to the following observation.

Observation 11. *A change in the visibility graph cannot increase the current configuration’s stretch.*




Fig. 6. The configuration polygon and its vertices of a configuration c and two examples for the inner angles $\alpha_\iota(t)$ (the image omits the time parameter to improve readability). Note that it is degenerated at several places, causing a single robot position to represent possibly multiple vertices (at most 6 by Lemma 14). The stretch is the length of the red boundary (where degenerated parts are counted twice).

5.3 Analysis of the MOVE-ON-BISECTOR Protocol

We are now ready to sketch the analysis of the MOVE-ON-BISECTOR protocol. The interested reader will find the full analysis in [5]. First we show, similar to the proof of Theorem 6 a lower bound on the rate at which the current configuration’s stretch decreases as long as the robots have not yet gathered (Lemmas 12 and 13). Afterward, we argue that a configuration’s stretch is linear in its diameter. Combining these easily yields the desired bound on the gathering time of MOVE-ON-BISECTOR (Theorem 15).

Lemma 12. *The stretch of a configuration at time t changes at a rate of*

$$\dot{s}(t) \leq -2 \sum_{\iota \in \mathcal{W}(t)} \cos(\alpha_\iota(t)/2). \tag{12}$$

Proof. Fix a time $t \geq 0$ and $\iota \in [q(t)]$. We consider how the distance $d(t) := \|w_\iota(t) - w_\iota(t)\|_2$ changes. Let $s_\iota(t)$ denote the speeds of vertex ι at time t .² Set $\beta_\iota(t) := \alpha_\iota(t)$ if $\iota \in \mathcal{W}(t)$ and $\beta_\iota(t) := 2\pi - \alpha_\iota$ if $\iota \notin \mathcal{W}(t)$. By Lemma 3, the distance changes at a rate of

$$\dot{d}(t) = -s_\iota(t) \cdot \cos(\beta_\iota(t)/2) - s_{\iota-1}(t) \cdot \cos(\beta_{\iota-1}(t)/2). \tag{13}$$

Note that the speed of any vertex $w_\iota(t)$ with $\iota \in \mathcal{W}(t)$ equals 1. Summing over all $\iota \in \{1, 2, \dots, q(t)\}$ we get that the current configuration’s stretch changes

² As in the proof of Theorem 6, we identify vertices with the robots positioned on them.

at a rate of $\dot{\mathfrak{s}}(t) = -2 \sum_{\iota \in [q(t)]} s_{\iota}(t) \cdot \cos(\beta_{\iota}(t)/2) \leq -2 \sum_{\iota \in \mathcal{W}} \cos(\beta_{\iota}(t)/2) = -2 \sum_{\iota \in \mathcal{W}} \cos(\alpha_{\iota}(t)/2)$. \square

Lemma 13. *The stretch of a configuration at time t decreases at a rate of at least 4.*

Proof. Since the configuration polygon is a (possibly degenerated) simple polygon, the sum of its inner angles $\alpha_{\iota}(t)$ is exactly $(q(t) - 2) \cdot \pi$. Define the angles $\beta_{\iota}(t) := \alpha_{\iota}(t)$ if $\iota \in \mathcal{W}(t)$ and $\beta_{\iota}(t) := 2\pi - \alpha_{\iota}(t)$ if $\iota \notin \mathcal{W}(t)$. Together with Lemmas 12 and 2(a) we get that the current configuration’s stretch decreases at a rate of at least

$$\begin{aligned} 2 \sum_{\iota \in \mathcal{W}(t)} \cos(\alpha_{\iota}(t)/2) &\geq 2 \sum_{\iota \in \mathcal{W}(t)} \left(1 - \frac{\alpha_{\iota}(t)}{\pi}\right) = 2 \sum_{\iota \in \mathcal{W}(t)} \left(1 - \frac{\alpha_{\iota}(t)}{\pi}\right) \\ + 2 \sum_{\iota \notin \mathcal{W}(t)} \left(1 - \frac{\pi}{\pi}\right) &\geq 2 \sum_{\iota \in [q(t)]} \left(1 - \frac{\alpha_{\iota}(t)}{\pi}\right) = 2q(t) - 2(q(t) - 2) = 4, \end{aligned} \tag{14}$$

finishing the proof. \square

Together, Observation 11 and Lemma 13 yield a bound of $O(\mathfrak{s})$ on the gathering time. To get our desired bound of $O(n)$, we use a result from [5] showing that $\mathfrak{s} \leq 6n$. This might seem trivial at first glance: After all, the stretch \mathfrak{s} corresponds to the length of a chain of actual robots that forms the configuration polygon. However, note that while the configuration polygon is simple, it might be degenerated, such that, for example, parts of it may form a line. Thus, the robots along the above-mentioned chain are not necessarily unique. However, using the basic geometry of the underlying visibility graph, one can show that no node appears more than 6 times in this chain.

Lemma 14 ([5, Lemma 5.7]). *For any configuration of n robots and stretch \mathfrak{s} , we have $\mathfrak{s} \leq 6n$.*

Combining Observation 11 and Lemmas 13 and 14, we immediately get the following result.

Theorem 15. *Consider an initial configuration of n robots and stretch \mathfrak{s} . In the worst-case, the MOVE-ON-BISECTOR protocol gathers the robots in time at most $\mathfrak{s}/4 \leq 3n/2$. This is asymptotically optimal for the gathering problem.*

6 Avoiding Collisions

The final section of this chapter gives a brief outlook of how one can avoid (early) collisions when gathering multiple point robots. Here, by collision we mean situations where two or more robots share the same position in the Euclidean plane. A configuration is called *collision-free* if no two robots share the same position. When the robots are gathered, we have a collision between all robots. We call this the *final collision*, which is not avoidable if we want to gather. Any collision between two or more robots before all robots are gathered is called an *early collision*. The question we study in this section is:

Are there continuous gathering protocols that, if started in a collision-free configuration, guarantee that no early collisions happen?

Such a gathering protocol is called *collision-free*.

Many natural gathering protocols, including the MOVE-ON-BISECTOR protocol, are prone to early collisions. In fact, in the analysis of both discrete and continuous setting, collisions are often seen as a success, as robots that collided behave identical (in synchronous and deterministic robot formation protocols). Thus, there can be at most $n - 1$ collisions, allowing us to use the number of collisions as a progress measure. However, from a practical standpoint and, collisions should be avoided as much as possible.

The rest of this section surveys recent results [12, 13] that made progress towards answering the above question.

6.1 A Candidate for an Almost Collision-Free Gathering Protocol

The GO-TO-THE-CENTER protocol is another simple and natural gathering protocol. Here, each robot moves with speed 1 towards the center of the *minimum enclosing circle* of all robot positions it currently sees.

In the discrete setting, this algorithm was introduced by Ando et al. [1] who showed that it indeed gathers all robots, but the authors provided no bound on the gathering time. A quadratic upper bound in this discrete setting was later shown by Degener et al. [4]. In the continuous setting, we can simply use our framework from Sect. 4: Indeed, it is easy to verify that GO-TO-THE-CENTER is a contracting protocol [13]. This allows us to apply Corollary 7 to get a quadratic bound on its gathering time.

Unfortunately, GO-TO-THE-CENTER is also prone to collisions. In fact, the analysis of its discrete variant [4] is partly based on collisions. However, as noted by Li et al. [12], one can slightly change the definition of GO-TO-THE-CENTER to get a promising candidate for an collision-free, continuous gathering protocol (or almost collision-free; see below).

GO-TO-THE-GABRIEL-CENTER. We can rephrase GO-TO-THE-CENTER as follows: each robot moves with speed 1 towards the center of the minimum enclosing circle of its (inclusive) neighborhood in the visibility graph. A variant of GO-TO-THE-CENTER – called GO-TO-THE-GABRIEL-CENTER and due to Li et al. [12] – is defined using the same phrasing but uses the so-called *Gabriel graph* [7] instead of the visibility graph:

Definition 16 (Gabriel Graph). *The Gabriel graph $GG_c = (VG_c, EG_c)$ of configuration c is a subgraph of the visibility graph. It has the same vertex set $VG_c = V_c = \{c_i \mid i \in [n]\}$. The edge set $EG_c \subseteq E_c$ consists of all edges $\{u, v\} \in E_c$ such that the interior of the smallest enclosing circle of u and v does not contain another robot’s position.*

See Fig. 7 for an example of how the Gabriel graph differs from the visibility graph.

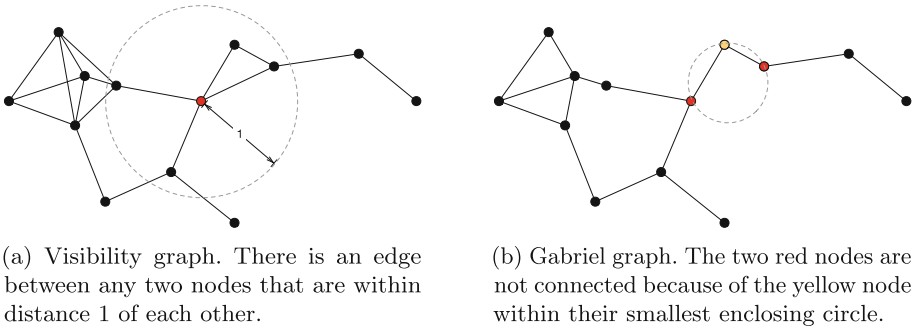


Fig. 7. Visibility graph vs. Gabriel graph.

With this definition, we can now formally define the GO-TO-THE-GABRIEL-CENTER protocol. Here, at each time t each robot i performs the following actions:

- (a) Robot i computes the minimum enclosing circle $\mathcal{C}_i(t)$ of all robots in its neighborhood (including i itself) of the Gabriel graph $GG_{c(t)}$. Let $T_i(t)$ denote the center of $\mathcal{C}_i(t)$ (robot i 's target point).
- (b) If $c_i(t)$ equals $T_i(t)$, robot i moves with $T_i(t)$.
- (c) If $c_i(t)$ is different from $T_i(t)$, robot i moves with speed 1 towards $T_i(t)$.

As with the GO-TO-THE-CENTER protocol, it is easy to see (and has been proved in [13]) that the GO-TO-THE-GABRIEL-CENTER protocol is contracting and, thus, gathers in quadratic time.


6.2 Collisions in the GO-TO-THE-GABRIEL-CENTER Protocol

Experimental results indicate that for “typical” initial configurations, the GO-TO-THE-GABRIEL-CENTER protocol causes no early collisions. In fact, in the one dimensional case (where all robots start on a line) this can be easily proved [12]. However, in the two dimensional case there are, in fact, some (quite symmetric) situations that exhibit early collisions. See Fig. 8 for an example. One of the central open questions left in [12] is whether the following conjecture is true.

Conjecture 17. *The set of initial configurations that lead to early collisions of the GO-TO-THE-GABRIEL-CENTER has Lebesgue measure 0.*

A less formal variant of this conjecture states that independent small random perturbations of the robots’ initial positions ensure that, with probability 1, there will be no early collisions. Note that even for the special case of $n = 4$ robots, a proof of Conjecture 17 seems challenging.

Even if Conjecture 17 is true, it remains open whether a distributed continuous gathering protocol can achieve truly collisionless gathering. There is currently



(a) By the configuration's symmetry, it is easy to see that the robots at the top and bottom will collide early in the center.

(b) A slight perturbation that destroys the symmetry is enough to ensure that no early collision occurs.

Fig. 8. Early collisions in the GO-TO-THE-GABRIEL-CENTER protocol.

only one continuous gathering protocol [13] that completely avoids early collisions, but it requires a considerably more complex robot model. Namely, robots must be non-oblivious (they have one memory bit), chiral (they share a common left/right orientation), and luminous (the contents of a robot's memory bit is visible to other robots). Additionally, this protocol requires quadratic time to gather all robots, so it is much slower than the MOVE-ON-BISECTOR protocol from Sect. 5. So even using a more complex robot model like the one above, it would be interesting to find a collision-free continuous gathering protocol that has linear gathering time.

References

1. Ando, H., Suzuki, I., Yamashita, M.: Formation and agreement problems for synchronous mobile robots with limited visibility. In: Proceedings of 10th International Symposium on Intelligent Control (ISIC), pp. 453–460, August 1995. <https://doi.org/10.1109/ISIC.1995.525098>
2. Cieliebak, M., Flocchini, P., Prencipe, G., Santoro, N.: Solving the robots gathering problem. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 1181–1196. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-45061-0_90
3. 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>
4. Degener, B., Kempkes, B., Langner, T., Meyer auf der Heide, F., Pietrzyk, P., Wattenhofer, R.: A tight runtime bound for synchronous gathering of autonomous robots with limited visibility. In: Rajaraman R., Meyer auf der Heide, F., (eds.) SPAA 2011: Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures, San Jose, CA, USA, 4–6 June 2011 (Co-located with FCRC 2011), pp. 139–148. ACM (2011). <https://doi.org/10.1145/1989493.1989515>
5. Degener, B., Kempkes, B., Kempkes, P., Meyer auf der Heide, F.: Linear and competitive strategies for continuous robot formation problems. *TOPC* **2**(1), 2:1–2:18 (2015). <https://doi.org/10.1145/2742341>

6. Dieudonné, Y., Petit, F.: Self-stabilizing deterministic gathering. In: Dolev, S. (ed.) ALGOSENSORS 2009. LNCS, vol. 5804, pp. 230–241. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-05434-1_23
7. Gabriel, K.R., Sokal, R.R.: A new statistical approach to geographic variation analysis. *Syst. Biol.* **18**(3), 259–278 (1969). <https://doi.org/10.2307/2412323>
8. Gordon, N., Wagner, I.A., Bruckstein, A.M.: Gathering multiple robotic a(ge)nts with limited sensing capabilities. In: Dorigo, M., Birattari, M., Blum, C., Gambardella, L.M., Mondada, F., Stützle, T. (eds.) ANTS 2004. LNCS, vol. 3172, pp. 142–153. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28646-2_13
9. Izumi, T., Katayama, Y., Inuzuka, N., Wada, K.: Gathering autonomous mobile robots with dynamic compasses: an optimal result. In: Pelc, A. (ed.) DISC 2007. LNCS, vol. 4731, pp. 298–312. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75142-7_24
10. Katayama, Y., Tomida, Y., Imazu, H., Inuzuka, N., Wada, K.: Dynamic compass models and gathering algorithms for autonomous mobile robots. In: Prencipe, G., Zaks, S. (eds.) SIROCCO 2007. LNCS, vol. 4474, pp. 274–288. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-72951-8_22
11. Katreniak, B.: Convergence with limited visibility by asynchronous mobile robots. In: Kosowski, A., Yamashita, M. (eds.) SIROCCO 2011. LNCS, vol. 6796, pp. 125–137. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22212-2_12
12. Li, S., Meyer auf der Heide, F., Podlipyan, P.: The impact of the gabriel subgraph of the visibility graph on the gathering of mobile autonomous robots. In: Chrobak, M., Fernández Anta, A., Gaşieniec, L., Klasing, R. (eds.) ALGOSENSORS 2016. LNCS, vol. 10050, pp. 62–79. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-53058-1_5
13. Li, S., Markarian, C., Meyer auf der Heide, F., Podlipyan, P.: A continuous strategy for collisionless gathering. In: Fernández Anta, A., Jurdzinski, T., Mosteiro, M.A., Zhang, Y. (eds.) ALGOSENSORS 2017. LNCS, vol. 10718, pp. 182–197. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-72751-6_14
14. Nguyen, H.G., Pezeshkian, N., Raymond, S.M., Gupta, A., Spector, J.M.: Autonomous communication relays for tactical robots. In: Proceedings of the 11th International Conference on Advanced Robotics (ICAR), pp. 35–40 (2003)
15. Olver, F.W., Lozier, D.W., Boisvert, R.F., Clark, C.W.: NIST Handbook of Mathematical Functions, 1st edn. Cambridge University Press, New York (2010)
16. Prencipe, G.: Impossibility of gathering by a set of autonomous mobile robots. *Theor. Comput. Sci.* **384**(2–3), 222–231 (2007). <https://doi.org/10.1016/j.tcs.2007.04.023>
17. Souissi, S., Défago, X., Yamashita, M.: Gathering asynchronous mobile robots with inaccurate compasses. In: Shvartsman, M.M.A.A. (ed.) OPODIS 2006. LNCS, vol. 4305, pp. 333–349. Springer, Heidelberg (2006). https://doi.org/10.1007/11945529_24
18. Watton, A., Kydon, D.W.: Analytical aspects of the n -bug problem. *Am. J. Phys.* **37**(2), 220–221 (1969). <https://doi.org/10.1119/1.1975458>