

Local Gathering of Mobile Robots in Three Dimensions (Full Version) *

Michael Braun, Jannik Castenow, and Friedhelm Meyer auf der Heide

Heinz Nixdorf Institute and Department of Computer Science
Paderborn University, Fürstenallee 11, 33102 Paderborn, Germany

{braunm, janniksu, fmadh}@mail.upb.de

Abstract

In this work, we initiate the research about the GATHERING problem for robots with limited viewing range in the three-dimensional Euclidean space. In the GATHERING problem, a set of initially scattered robots is required to gather at the same position. The robots' capabilities are very restricted – they do not agree on any coordinate system or compass, have a limited viewing range, have no memory of the past and cannot communicate.

We study the problem in two different time models, in $\mathcal{F}\text{SYNC}$ (fully synchronized discrete rounds) and the continuous time model. For $\mathcal{F}\text{SYNC}$, we introduce the 3D-GO-TO-THE-CENTER-strategy and prove a runtime of $\Theta(n^2)$ that matches the currently best runtime bound for the same model in the Euclidean plane [SPAA'11].

Our main result is the generalization of contracting strategies (continuous time model) from [Algosensors'17] to the three-dimensional case. In contracting strategies, every robot that is located on the global convex hull of all robots' positions moves with full speed towards the inside of the convex hull. We prove a runtime bound of $\mathcal{O}(\Delta \cdot n^{3/2})$ for *any* three-dimensional contracting strategy, where Δ denotes the diameter of the initial configuration. This comes up to a factor of \sqrt{n} close to the lower bound of $\Omega(\Delta \cdot n)$ which is already true in two dimensions.

In general, it might be hard for robots with limited viewing range to decide whether they are located on the global convex hull and which movement maintains the connectivity of the swarm, rendering the design of *concrete* contracting strategies a challenging task. We prove that the continuous variant of 3D-GO-TO-THE-CENTER is contracting and keeps the swarm connected. Moreover, we give a simple design criterion for three-dimensional contracting strategies that maintains the connectivity of the swarm and introduce an exemplary strategy based on this criterion.

1 Introduction

We study a scenario where a distributed system of mobile entities (called *robots*) is supposed to establish a certain formation, also denoted as a *pattern*. The robots are scattered in a d -dimensional Euclidean space (usually the Euclidean plane) and have to coordinate their movements in a distributed manner to reach the desired formation. The robots' capabilities depend on the exact model and formation problem but are typically very restricted. Usually, the robots do not agree on a common coordinate system or compass, cannot communicate with each other and have only limited

*This paper is a full version of the respective paper presented at SIROCCO 2020.

sensing capabilities. One extensively studied coordination problem is the PATTERN FORMATION problem, dealing with questions such as: Which patterns are generally formable by a set of robots? Which capabilities do the robots need? Given a specific pattern, for which initial configurations is this pattern formable? Interestingly, it has been proven that there are only *two* patterns that might be formable starting in an arbitrary input configuration. These are the patterns POINT and UNIFORM CIRCLE. Forming the pattern POINT is known under a more common name – the GATHERING problem, which studies the task of gathering a set of robots on the same position. Both of these problems have been extensively studied under several different assumptions, involving the viewing range (local or global), the synchronization (synchronous or asynchronous activation), the extent (robots can or cannot occupy the same position) or the opacity of robots, to name only a few. However, most of these models have in common that the robots operate in the two-dimensional Euclidean plane. A natural extension would be to consider the three-dimensional Euclidean space, where the robots have the ability to fly, such as drones, or to move underwater. Existing results about robots in the three-dimensional Euclidean space are very scarce, rely on strong assumptions (such as axis agreement) and do not consider any runtime analyses of the proposed strategies. Our work initiates the study of GATHERING of robots in three-dimensions, in one of the weakest possible models – robots do not agree on any coordinate system or compass, are *oblivious* (have no memory of the past) and have only a local view.

1.1 Model & Time Notions

We consider a set \mathcal{R} of n robots r_1, \dots, r_n , each of which occupies a single point in \mathbb{R}^3 at each time. As such, robots can neither block each other's views nor paths, and multiple robots are allowed to occupy the same position at the same time. The position of robot r_i at time t is denoted by $p_i(t)$. The positions of all robots at time t , $\mathcal{P}_t = (p_1(t), \dots, p_n(t))$ are collectively called the *configuration* at time t . The Euclidean distance between points $x, y \in \mathbb{R}^3$ is denoted as $d(x, y)$. For a subset of the three-dimensional Euclidean space $\mathcal{P} \subseteq \mathbb{R}^3$, $d(x, \mathcal{P})$ is used as a shorthand for $\min_{y \in \mathcal{P}} d(x, y)$.

The overall abilities of the robots are rather limited: They are not allowed to communicate with each other, they are *identical* (they cannot be distinguished) and are *oblivious*, meaning they have no memory of the past. Furthermore they do not share a common coordinate system or orientation. Robots are only able to observe the space around them within a limited viewing range of 1, i.e. a robot r_i can see the position of another robot r_j if and only if $d(p_i(t), p_j(t)) \leq 1$. Two robots r_i and r_j with $d(p_i(t), p_j(t)) \leq 1$ are also called *neighbors*. The set of all neighbors of r_i at time t is called the *neighborhood* of r_i and is denoted as $\mathcal{R}_i(t)$. This limited viewing range can also be considered to induce a unit ball graph $\text{UBG}_t = (\mathcal{R}, E_t)$ at time t , whose nodes consist of the robots and where the set of edges E_t contains an edge $\{r_i, r_j\}$ if and only if $d(p_i(t), p_j(t)) \leq 1$. This graph is also called the *visibility graph* at time t . Note that the UBG is a generalization of the two-dimensional *unit disk graph* (UDG) to three dimensions.

Starting from a configuration of n robots in the three-dimensional Euclidean space that is connected at time 0, i.e. UBG_0 is connected, the goal is to gather all robots in one point. This problem will be referred to as the (three-dimensional) GATHERING problem. Note that the eventual gathering point is not predefined and can instead be chosen by the robots at runtime. This also imposes a subgoal during the execution of any algorithm that solves this problem: It has to be ensured that UBG_t remains connected. Otherwise, the limited viewing range of the robots, combined with the fact that they do not share coordinate systems, makes it impossible for any deterministic algorithm to restore connectivity and the robots can no longer converge to the same point [1].

Throughout this work, we consider two different notions of time: The fully synchronous \mathcal{FSYNC} model and the continuous time model.

$\mathcal{F}\text{sync}$: In $\mathcal{F}\text{sync}$, all robots operate in fully synchronous **Look-Compute-Move** (LCM) cycles. In the **Look** phase, a robot r_i observes its environment, detects the set of all visible robots $\mathcal{R}_i(t)$ and stores a snapshot in its local memory. Based on this snapshot, r_i computes a target point in the **Compute** phase. Finally, in the **Move** phase, r_i moves to that target point. The execution of a single LCM cycle is also denoted as one *round*.

Continuous Time Model: Generally, the continuous time model can be seen as a continuous variant of $\mathcal{F}\text{sync}$, in which robots only move an infinitesimal small distance towards their target points [7]. At every point in time, the movement of each robot r_i can be expressed by a *velocity vector* $\vec{v}_i(t)$ with $0 \leq \|\vec{v}_i(t)\| \leq 1$, i.e. the maximal speed of a robot is bounded by 1. In contrast to $\mathcal{F}\text{sync}$, the function $p_i: \mathbb{R}_{>0} \rightarrow \mathbb{R}^3$, representing the position of r_i at time t , is a continuous function and also called the *trajectory* of r_i . Although the trajectories are continuous, they are not necessarily differentiable because robots are able to change their speed and direction non-continuously. However, natural movement strategies have (right) differentiable trajectories. Thus, the velocity vector of a robot $\vec{v}_i: \mathbb{R}_{>0} \rightarrow \mathbb{R}^3$ can be seen as the (right) derivative of p_i .

1.2 Our Contribution

The contribution of this paper is twofold. We consider the fully synchronous $\mathcal{F}\text{sync}$ model and the continuous time model. For $\mathcal{F}\text{sync}$, we introduce the strategy 3D-GO-TO-THE-CENTER (3D-GTC), which is the three-dimensional generalization of GO-TO-THE-CENTER (GTC), invented for robots operating in the Euclidean plane [1]. The main idea of 3D-GTC is that robots move towards the center of the smallest enclosing sphere of all robots within their viewing radius, while ensuring that the configuration stays connected. We prove a runtime bound of $\Theta(n^2)$ for 3D-GTC which matches the runtime of the two-dimensional GTC strategy.

For the continuous time model, we generalize the class of contracting strategies [9] to three dimensions. In contracting strategies, every robot that lies on the convex hull of all robots' positions moves always with speed 1 into a direction that points inside or on the boundary of the convex hull. We prove that every (three-dimensional) contracting gathering strategy gathers all robots on a single point in time at most $\mathcal{O}(\Delta \cdot n^{3/2})$, where Δ denotes the (geometric) diameter of the initial configuration, i.e. the maximum Euclidean distance between any pair of robots. This runtime bound differs from the runtime bound for two-dimensional contracting strategies by a factor of \sqrt{n} . The lower bound is $\Omega(\Delta \cdot n)$ and already holds for the two-dimensional case [9]. The main open question is whether $\mathcal{O}(\Delta \cdot n^{3/2})$ is tight or can be improved to $\mathcal{O}(\Delta \cdot n)$.

Note that a contracting strategy is not necessarily local. Therefore, we finally present two local, contracting strategies. Our first example is the continuous variant of 3D-GTC, called CONT-3D-GTC. We prove that the strategy is contracting and thus gathers the robots in time $\mathcal{O}(\Delta \cdot n^{3/2})$. In addition, we present the class of *tangential-normal* strategies. These strategies are local and maintain connectivity. As an example for a strategy that is both tangential-normal and contracting, we introduce the MOVE-ON-ANGLE-MINIMIZER strategy.

1.3 Related Work

In this overview over related work, we focus on the GATHERING problem for synchronized robots with local visibility in the Euclidean plane. Beyond that, we give a summary about research concerning robot coordination problems in the three-dimensional Euclidean space. For other models and coordination problems, which involve, among others, less synchronized schedulers or robots with a global view, we refer the reader to the recent survey [6].

Ando, Suzuki and Yamashita introduced the GTC-strategy for fully synchronous robots with local view [1]. In GTC, every robot moves in every round towards the center of the smallest enclosing circle of all robots within its viewing range while ensuring that the swarm remains connected. Ando et al. could prove that GTC solves the GATHERING-problem in finite time. Later on, Degener et al. could prove a tight runtime bound of $\Theta(n^2)$ for GTC [4]. By now, this is the best known runtime bound for a strategy that solves GATHERING of robots with local visibility and without agreement on any coordinate system or compass in \mathcal{FSYNC} .

Faster runtimes could so far only be obtained under different assumptions – for example by introducing one-axis agreement or changing the time model. Poudel and Sharma proved that it is possible to gather a swarm of robots with local view in time $\mathcal{O}(\Delta)$, where Δ denotes the diameter of the initial configuration [11]. The main assumption for their strategy is that the robots agree on one axis of their coordinate systems.

The second time model we consider in this paper is the continuous time model, introduced by Gordon et al. [7]. In this time model robots do not operate in synchronized rounds but continuously observe their environment and move while having a bounded maximal speed. Gordon et al. propose a gathering strategy for the continuous time model. In their strategy, all robots that locally assume that they are located on the global convex hull move with maximal speed along the bisector formed by vectors to their neighbors along the global convex hull. This strategy has later been called MOVE-ON-BISECTOR by Degener et al. They could also prove runtime of $\Theta(n)$ [3].

The main result of this paper is based on a more general view on continuous GATHERING strategies in the Euclidean plane – the class of *contracting* strategies in which all robots that are located on the global convex hull of all robots move with maximal speed into a direction that points inside of the convex hull [9]. Li et al. could prove a runtime of $\mathcal{O}(\Delta \cdot n)$ for *any* contracting strategy. Note that MOVE-ON-BISECTOR is also a contracting strategy but has a significantly faster runtime than $\mathcal{O}(\Delta \cdot n)$. However, there are contracting strategies with a runtime of $\Omega(\Delta \cdot n)$ [9].

In the three-dimensional Euclidean space there is so far, to the best of our knowledge, no strategy known that solves GATHERING of robots with limited viewing range. More generally, literature about robots operating in three-dimensional spaces is very scarce. We summarize the literature briefly. In [2] the authors show that gathering of robots in the three-dimensional Euclidean space is possible – under the assumptions that robot have a global view but are not transparent and that the robots agree on one axis of their coordinate systems. Tomaita et al. study a different problem – the plane formation problem [12]. In the plane formation problem, the goal is that eventually all robots are located on the same plane, while ensuring that no two robots occupy the same position. The authors show that this problem is not solvable for every initial configuration, give a characterization of all start configurations for which the problem is solvable and introduce an algorithm that solves the problem for the latter set of configurations. Yamauchi, Uehara and Yamashita generalize this result further and study the more general PATTERN FORMATION for synchronized robots in the three-dimensional Euclidean space [13]. They characterize the set of all patterns that might be formable depending on symmetries of the initial configuration.

2 Gathering in \mathcal{Fsync}

In this section, the three-dimensional GATHERING problem will be studied under the \mathcal{FSYNC} model. The results can be considered as a generalization of those obtained by Degener et al. [4] for the two-dimensional setting. It will be shown that a generalization of GTC by Ando et al. [1] solves the gathering problem in three dimensions in $\Theta(n^2)$ rounds.

2.1 3d-Go-To-The-Center

Algorithm 1 3D-GO-TO-THE-CENTER (3D-GTC)

- 1: $\mathcal{R}_i(t) := \{\text{positions of robots visible from } r_i, \text{ including } r_i \text{ at time } t\}$
 - 2: $\mathcal{S}_i(t) := \text{smallest enclosing sphere of } \mathcal{R}_i(t)$
 - 3: $c_i(t) := \text{center of } \mathcal{S}_i(t)$ ▷ target point
 - 4: **for all** $r_j \in \mathcal{R}_i(t)$ **do** ▷ Maintain connectivity
 - 5: $m_j := \text{midpoint between } p_i(t) \text{ and } p_j(t)$
 - 6: $\mathcal{B}_j(t) := \text{ball with radius } \frac{1}{2} \text{ and center } m_j$
 - 7: $\ell_j := \text{maximum distance } r_i \text{ can move towards } c_i(t) \text{ without leaving } \mathcal{B}_j(t)$
 - 8: $L_i := \min_{r_j \in \mathcal{R}_i(t)} \ell_j$
 - 9: Move towards $c_i(t)$ for a distance of L_i
-

The strategy 3D-GO-TO-THE-CENTER (3D-GTC) is a generalization of GO-TO-THE-CENTER to the three-dimensional Euclidean space and is summarized in Algorithm 1. A key component is the computation of a smallest enclosing sphere (SES) of a set of points \mathcal{P} . This is a sphere of minimal radius that contains all points in \mathcal{P} with the following properties:

Proposition 1. [5] *Let \mathcal{S} be the smallest enclosing d -sphere (SES) of a point set $\mathcal{P} \subset \mathbb{R}^d$. Then the center c of \mathcal{S} is a convex combination of at most $d + 1$ points in \mathcal{P} that lie on the surface of \mathcal{S} . Especially,*

1. c lies in \mathcal{S}
2. c minimizes the maximum distance to the points in \mathcal{P} .

Intuitively, 3D-GTC works by attempting to locally move robots closer together. This is achieved by letting each robot r_i compute the SES of its neighborhood $\mathcal{R}_i(t)$ and then moving towards its center $c_i(t)$. Additionally, the strategy follows the subgoal of maintaining connectivity of UBG_{t+1} . This is achieved by limiting the distance a robot r_i moves towards its target $c_i(t)$, such that for any of its neighbors r_j , it stays within a distance of $\frac{1}{2}$ of the midpoint between the positions of r_i and r_j at time t . Thus, if both r_i and r_j perform this strategy, the distance between their positions at the start of the next round $t + 1$ is at most 1, maintaining visibility. By the argumentation above, the following Lemma holds.

Lemma 2. *If UBG_0 is connected, UBG_t remains connected for all $t \geq 0$.*

Overall, the only difference to the original GTC strategy for two dimensions lies in the computation of a smallest enclosing sphere in the 3D case over a smallest enclosing circle in the 2D case. In fact, if the three-dimensional version is applied to a configuration of robots that is coplanar with respect to some plane h , it acts just as if the robots' positions were projected to h and the two-dimensional version was applied to the resulting two-dimensional subspace. This is a result of the fact that computing a SES of a set of coplanar points is equivalent to computing a smallest enclosing circle instead.

From this observation, we can immediately conclude that the lower bound on the runtime of the two-dimensional version of the strategy shown by Degener et al. [4] also applies to the three-dimensional case by simply embedding the two-dimensional worst-case start configuration within three-dimensional space: In the configuration, n robots are positioned on a circle such that the distance between two neighbors is 1. This causes the robots to only take small steps of size $\mathcal{O}(1/n)$ towards the center of the circle, leading to a gathering time of $\Omega(n^2)$.

Theorem 3. *There is a start configuration such that 3D-GTC takes $\Omega(n^2)$ rounds to gather the robots in one point.*

With a generalization of the analysis of [4], we can also prove an upper runtime bound of $\mathcal{O}(n^2)$. Due to space constraints, the analysis is moved to Appendix A.

Theorem 4. *Given n robots in a connected starting configuration $\mathcal{P} \in \mathbb{R}^3$ in the Euclidean space, 3D-GTC gathers the robots in $\mathcal{O}(n^2)$ rounds.*

The combination of both theorems yields a tight runtime of $\Theta(n^2)$.

3 Continuous Gathering

Now, we consider the GATHERING problem within the continuous time model. For the Euclidean plane, Li et al. [9] introduced the class of *contracting strategies*. This definition can also be applied to three dimensions: Let CH_t denote the closed convex hull of the robots' configuration \mathcal{P}_t at time t and let Corn_t denote the vertices of CH_t . The class of contracting strategies can be defined as follows:

Definition 1. *In the continuous time model, a movement strategy for n robots is called contracting if for every time t such that the cardinality of Corn_t is strictly greater than 1, every robot in Corn_t moves with speed 1 in a direction that points to CH_t .*

The main idea of our analysis is to project the three-dimensional configuration (including the velocity vectors) to a two-dimensional plane. The projected robots then perform something similar to a contracting strategy where they move towards the inside of the projected convex hull with varying speeds. However, when looking at only a single projection plane, some velocity vectors might even have a length of 0 in the projection at some points in time (in case the projection plane is chosen orthogonal to the velocity vector). Thus, the analysis of Li et al. cannot be directly applied to the projection as this analysis assumes that all robots on the convex hull move with speed 1 towards the inside. Instead, we analyze not only one but all possible (meaningfully different) projections, since – intuitively – for a majority of all possible projection planes, the projected length of a velocity vector must be larger than a constant ε .

3.1 Preliminaries

The following lemma is a useful tool for the analysis of continuous strategies stating how the distance between two robots changes over time.

Lemma 5 ([8]). *Consider two robots r_i and r_j with differentiable trajectories at time t . Their distance $d(p_i(t), p_j(t))$ at time t changes with speed*

$$d'(p_i(t), p_j(t)) = -(\|v_i(t)\| \cdot \cos \beta_{i,j}(t) + \|v_j(t)\| \cdot \cos \beta_{j,i}(t)),$$

where $\beta_{i,j}(t)$ is the angle between $\vec{v}_i(t)$ and the line segment $\overline{p_i(t)p_j(t)}$.

The main tool for the analysis of contracting strategies in the three-dimensional Euclidean space are projections of the robots' configuration onto a two-dimensional plane. Let $h(\vec{x})$ be the plane through the origin with normal vector \vec{x} and let $\Pi_{\vec{x}}$ denote the orthogonal projection onto $h(\vec{x})$. Now, given a configuration \mathcal{P} of n robots, consider their projection $\hat{\mathcal{P}}^{(\vec{x})} = \{\Pi_{\vec{x}}p_i(t) \mid p_i(t) \in \mathcal{P}\}$

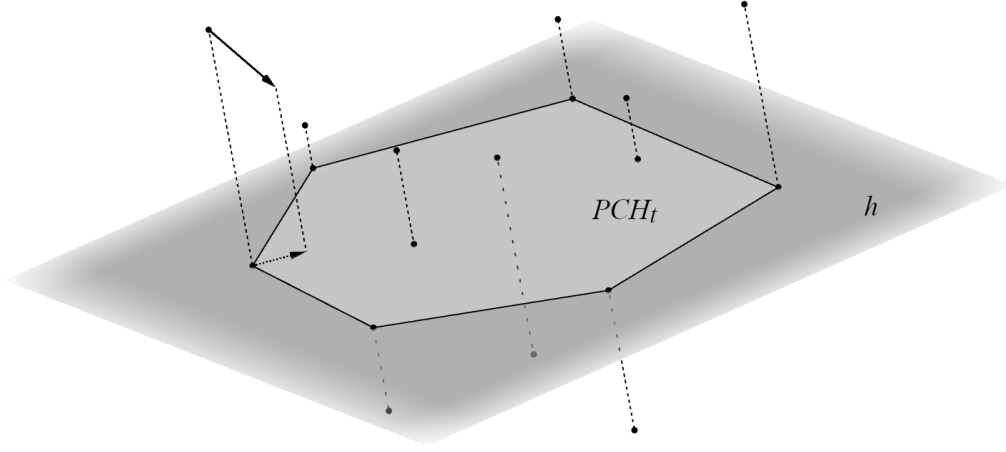


Figure 1: A configuration of robots being projected onto a plane $h(\vec{x})$. The mapping of the orthogonal projection $\Pi_{\vec{x}}$ is illustrated by dashed lines and the projected convex hull $PCH_t(\vec{x})$ is shown in light grey. One of the robots' movement vectors as well as its projection are represented by arrows.

onto $h(\vec{x})$ along with the projections of their movement vectors $\hat{v}_i^{(\vec{x})}(t) = \Pi_{\vec{x}}\vec{v}_i(t)$. Furthermore, denote the convex hull of $\hat{\mathcal{P}}$ as $PCH_t(\vec{x})$. See also Fig. 3.1.

If the robots perform a contracting strategy in the three-dimensional space, their projections also move towards the inside of the projected convex hull $PCH_t(\vec{x})$ since $\Pi_{\vec{x}}$ is a linear transformation and therefore preserves convexity. However, the lengths of the projected movement vectors $\hat{v}_i(t)$ are going to be smaller than 1 in general. For a given projection onto a plane $h(\vec{x})$, the minimum length of the $\hat{v}_i^{(\vec{x})}(t)$ will be called the *projected speed* and is denoted by $\varepsilon_{\vec{x}} = \min_{r_i \in \mathcal{R}} \|\Pi_{\vec{x}}\vec{v}_i(t)\|$. Note that $\varepsilon_{\vec{x}}$ can even be 0 in case $h(\vec{x})$ is orthogonal to any velocity vector. The following notion of the *length* of $PCH_t(\vec{x})$ will be used as a part of a progress measure for three-dimensional contracting strategies:

Definition 2. (*Length*) Let $m_1(t), m_2(t), \dots, m_{k(t)}(t)$ be the vertices of $PCH_t(\vec{x})$ (ordered counter-clockwise), where $k(t)$ is the number of vertices at time t . The length $\ell(t, \vec{x})$ of $PCH_t(\vec{x})$ is defined as the sum of its edge lengths: $\ell(t, \vec{x}) = \sum_{i=1}^{k(t)} d(m_i(t), m_{i-1}(t))$, where $m_0 := m_{k(t)}(t)$.

Note that if the diameter of the starting configuration was Δ , the length of a given projection can be at most $\pi\Delta$ (if it approximates a circle). Furthermore, if $\ell(t, \vec{x}) = 0$, then the robots have either gathered in the original three-dimensional space or have formed a line that is parallel to \vec{x} . In the latter case it only takes further time of at most $O(\Delta)$ for the robots to gather, as those robots that form the endpoints of the line have no choice but to move towards each other. The following Lemma provides a statement about how the length changes over time.

Lemma 6. For time t , let $h(\vec{x})$ be a plane with projected speed $\varepsilon_{\vec{x}}$, such that $\ell(t, \vec{x}) > 0$ and no two robots with different positions in \mathbb{R}^3 get projected onto the same point on $h(\vec{x})$. Then $\ell'(t, \vec{x}) \leq -\frac{8\varepsilon_{\vec{x}}}{n}$.

Proof. Because $\Pi_{\vec{x}}$ is a linear transformation, each of the $m_i(t)$ (corners of $PCH_t(\vec{x})$) must also be the projection of one of the vertices of the original, three-dimensional convex hull CH_t . Therefore,

they possess velocity vectors that point towards the inside of CH_t by the definition of a contracting strategy. Now consider the projections of these velocity vectors onto $h(\vec{x})$: Let $\hat{v}_i^{(\vec{x})}(t) := \Pi_{\vec{x}} \vec{v}_i(t)$. By assumption, we have $\|\hat{v}_i^{(\vec{x})}(t)\| \geq \varepsilon_{\vec{x}}$. Using this, it is now possible to bound $\ell'(t, \vec{x})$: Let $\alpha_i(t)$ be the internal angle of $\text{PCH}_t(\vec{x})$ at $m_i(t)$.

Note that in general, it may happen that two corner robots of CH_t got projected onto the same point on $h(\vec{x})$ for some \vec{x} . By one of the assumptions of the lemma, this is not true. Therefore, we know that each corner $m_i(t)$ of $\text{PCH}_t(\vec{x})$ contains only a single robot. This means that each $\alpha_i(t)$ is split into two parts, $\hat{\beta}_{i,\iota-1}(t)$ and $\hat{\beta}_{\iota-1,i}(t)$ by $m_i(t)$'s velocity vector $\hat{v}_i^{(\vec{x})}(t)$, such that $\alpha_i(t) = \hat{\beta}_{i,\iota-1}(t) + \hat{\beta}_{\iota-1,i}(t)$. Using Lemma 5 and Proposition ??, the derivative of $\ell(t)$ can now be bounded as follows: Recall that $\ell'(t, \vec{x}) = \sum_{\iota=1}^{k(t)} d'(m_\iota(t), m_{\iota-1}(t))$:

$$\ell'(t, \vec{x}) = \sum_{\iota=1}^{k(t)} d'(m_\iota(t), m_{\iota-1}(t)) \quad (1)$$

$$= \sum_{\iota=1}^{k(t)} -(\|\hat{v}_\iota^{(\vec{x})}(t)\| \cos \hat{\beta}_{\iota,\iota-1}(t) + \|\hat{v}_{\iota-1}^{(\vec{x})}(t)\| \cos \hat{\beta}_{\iota-1,\iota}(t)) \quad (2)$$

$$\leq -\varepsilon_{\vec{x}} \sum_{\iota=1}^{k(t)} \cos \hat{\beta}_{\iota,\iota-1}(t) + \cos \hat{\beta}_{\iota-1,\iota}(t) \quad (3)$$

$$= -\varepsilon_{\vec{x}} \sum_{\iota=1}^{k(t)} \frac{2(\alpha_\iota(t) - \pi)^2}{\pi^2} \quad (4)$$

$$= -\frac{2\varepsilon_{\vec{x}}}{\pi^2} \sum_{\iota=1}^{k(t)} (\alpha_\iota(t) - \pi)^2 \quad (5)$$

For Equation (4) observe that for $\vartheta \in [0, 1]$ and $\alpha \in [0, \pi]$, it holds that $\cos(\alpha\vartheta) + \cos(\alpha(1-\vartheta)) \geq \frac{2(\alpha-\pi)^2}{\pi^2}$ [10]. Now, the Cauchy-Schwarz inequality along with the fact that the sum of the inner angles of a convex polygon with k corners is $(k-2) \cdot \pi$.

$$\begin{aligned} \ell'(t, \vec{x}) &\leq -\frac{2\varepsilon_{\vec{x}}}{k(t) \cdot \pi^2} \cdot \left(\sum_{\iota=1}^{k(t)} (\alpha_\iota(t) - \pi) \right)^2 = -\frac{2\varepsilon_{\vec{x}}}{k(t) \cdot \pi^2} \cdot ((k(t)-2) \cdot \pi - k(t) \cdot \pi)^2 \\ &= -\frac{8\varepsilon_{\vec{x}}}{k(t)} \leq -\frac{8\varepsilon_{\vec{x}}}{n} \end{aligned}$$

This concludes the proof. □

Note that this also means that $\ell(t, \vec{x})$ is monotonically decreasing over time.

3.2 Proof of the upper bound

The main idea of the analysis is to track the lengths $\ell(t, \vec{x})$ for all (meaningfully different) projection planes $h(\vec{x})$. Since the length of the normal vector does not matter, it is enough to consider only vectors \vec{x} of length 1. Additionally, a vector \vec{x} and its reflection about the origin $-\vec{x}$ describe the same plane. Therefore it is enough to consider those vectors that lie on the surface of a unit hemisphere U centered around the origin (w.l.o.g. the one above the XY -plane).

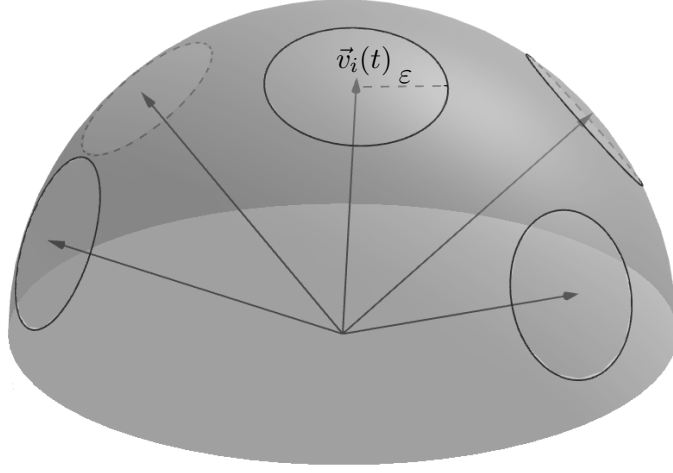


Figure 2: A figure illustrating how movement vectors block areas of the unit hemisphere U . Around each movement vector $\vec{v}_i(t)$, there is a spherical cap of radius ε . Each plane corresponding to a normal vector \vec{x} lying in one of those spherical caps is blocked.

The integral of the lengths $\ell(t, \vec{x})$ with respect to \vec{x} on the surface of U at time t can now be used as a measure to track the progress of a three-dimensional gathering strategy:

$$L(t) = \iint_U \ell(t, \vec{x}) dA$$

If $L(t) = 0$, the robots have gathered. If one of the $\ell(t, \vec{x})$ prematurely becomes 0, then the robots are collinear and gather in further time $O(\Delta)$.

Lemma 7. $L(0) \leq 2\pi^2 \Delta$.

Proof. Since $\ell(t, \vec{x}) \leq \pi\Delta$ (if $\text{PCH}_t(\vec{x})$ approximates a circle), we conclude

$$L(0) \leq \iint_U \pi\Delta dA = \pi\Delta \iint_U dA$$

The remaining integral part is a surface integral over a hemisphere. By observing that the surface area of a unit hemisphere is 2π , the lemma follows. \square \square

The goal of the proof is to show that there is at least a constant $(1 - \alpha)$ -fraction of projection planes $h(\vec{x})$ with projected speed at least ε for some constants α and ε . This can then be used to show that $L(t)$ decreases by a constant amount at each point in time using Lemma 6.

Now consider a projection plane $h(\vec{x})$. If this plane has projected speed smaller than ε at time t , then there is a movement vector $\vec{v}_i(t)$, such that $\angle(\vec{x}, \vec{v}_i(t)) < \sin^{-1} \varepsilon$. We say that $\vec{v}_i(t)$ *blocks* $h(\vec{x})$. Conversely, given a $\vec{v}_i(t)$, we can determine the set of all the $h(\vec{x})$ that are blocked by this $\vec{v}_i(t)$:

Lemma 8. *At time t , the movement vector $\vec{v}_i(t)$ blocks vectors from an area of $2\pi(1 - \sqrt{1 - \varepsilon^2})$ on U from reaching projected speed ε .*

Proof. W.l.o.g. it can be assumed that $\vec{v}_i(t)$ has a positive z -component, i.e. lies on U . Otherwise it can be reflected about the origin and it will still affect the exact same planes.

Now consider the spherical cap of U with base radius ε and apex $\vec{v}_i(t)$ and let C be its curved surface (see Fig. 3.2 for an illustration). For all vectors $\vec{x} \in C$, $h(\vec{x})$ is blocked from reaching projected speed ε . The area of C can be computed by $A_C = 2\pi r^2(1 - \cos \theta) = 2\pi(1 - \cos(\sin^{-1} \varepsilon)) = 2\pi(1 - \sqrt{1 - \varepsilon^2})$ \square \square

Since there are n robots, the area blocked by their movement vectors is at most $n \cdot 2\pi(1 - \sqrt{1 - \varepsilon^2})$, whereas the total surface of U is 2π . If we want the movement vectors to block only an α -fraction of U 's surface, the ε can be chosen accordingly:

Lemma 9. *Let $0 \leq \alpha \leq 1$. Then for a minimum speed of $\varepsilon = \frac{\sqrt{2n\alpha - \alpha^2}}{n}$, there is at most an α -fraction of the surface of U that is blocked with respect to ε .*

Proof. U has a surface of 2π and the robots' movement vectors block an area of at most $n \cdot 2\pi(1 - \sqrt{1 - \varepsilon^2})$. We want to choose ε such that the following holds:

$$\alpha 2\pi = n \cdot 2\pi(1 - \sqrt{1 - \varepsilon^2}) \iff \varepsilon = \frac{\sqrt{2n\alpha - \alpha^2}}{n}$$

\square
 \square

Using this lemma, it is now possible to bound the decrease of the progress measure $L(t)$ for a given α :

Lemma 10. *For a time $t \geq 0$ such that $\ell(t, \vec{x}) > 0$ for all $\vec{x} \in U$ and $0 \leq \alpha \leq 1$, then $L'(t) \leq -16\pi \cdot (1 - \alpha) \cdot \frac{\sqrt{2n\alpha - \alpha^2}}{n^2}$.*

Proof. Choose $\varepsilon = \frac{\sqrt{2n\alpha - \alpha^2}}{n}$ according to Lemma 9, i.e. there is only at most an α -fraction of the surface of U that is blocked. Since Lemma 6 only applies to those \vec{x} for which no two robots get projected onto the same point, the \vec{x} for which this is the case still have to be considered. However, there is only a finite number $\binom{n}{2}$ of such vectors out of the uncountably many that form U and they are only singular points on U . Therefore, they can be ignored when considering the integral $L(t)$. By Lemma 6, there is an $(1 - \alpha)$ -fraction of vectors \vec{x} from the surface of U (which has size 2π) with $\ell'(t, \vec{x}) \leq -\frac{8\varepsilon}{n} = -8\frac{\sqrt{2n\alpha - \alpha^2}}{n^2}$. Using this, we can bound $L'(t)$:

$$\begin{aligned} L'(t) &= \frac{d}{dt} \left(\iint_U \ell(t, \vec{x}) dA \right) = \iint_U \ell'(t, \vec{x}) dA \\ &\leq (1 - \alpha) \cdot 2\pi \cdot -8\frac{\sqrt{2n\alpha - \alpha^2}}{n^2} = -16\pi \cdot (1 - \alpha) \cdot \frac{\sqrt{2n\alpha - \alpha^2}}{n^2} \end{aligned}$$

\square \square

By choosing the α appropriately, the main result can now be obtained:

Theorem 11. *A set of n robots controlled by a contracting strategy gathers in time $\mathcal{O}(\Delta \cdot n^{3/2})$ from an initial configuration with diameter Δ .*

Proof. By Lemma 7, we have $L(0) \leq 2\pi^2\Delta$. By Lemma 10, $L(t)$ decreases by at least $16\pi \cdot (1 - \alpha) \cdot \frac{\sqrt{2n\alpha - \alpha^2}}{n^2}$ for a given α as long as $\ell(t, \vec{x}) > 0$ for all $\vec{x} \in U$. However, if there is an $\vec{x} \in U$ with $\ell(t, \vec{x}) = 0$, then the robots are collinear along some line that is parallel to \vec{x} and take further time $O(\Delta)$ to gather.

Now choose $\alpha = \frac{1}{2}$ and consider an arbitrary time t such that $\ell(t, \vec{x}) > 0$ for all $\vec{x} \in U$. Then $L'(t) \leq -8\pi \frac{1}{n^{3/2}}$. Therefore it takes time at most $(2\pi^2\Delta)/(8\pi \frac{1}{n^{3/2}}) = \frac{\pi}{4}\Delta n^{3/2}$ until $L(t)$ is zero. This leads to a gathering time of $\mathcal{O}(\Delta \cdot n^{3/2}) + \mathcal{O}(\Delta) \in \mathcal{O}(\Delta \cdot n^{3/2})$. \square \square

3.3 Continuous-3d-Go-To-The-Center

Next, a continuous version of 3D-GTC which was already presented for the discrete time setting, will be considered as a concrete example of a contracting strategy. The two-dimensional version of this strategy was adapted for continuous time by Li et al. [10]. Compared to the discrete time version, no additional measures have to be taken to preserve connectivity, as it can be shown that this happens naturally in the continuous case. The strategy is summarized in Algorithm 2.

Algorithm 2 CONTINUOUS-3D-GO-TO-THE-CENTER (CONT-3D-GTC)

- 1: $\mathcal{R}_i(t) := \{\text{positions of robots visible from } r_i, \text{ including } r_i \text{ at time } t\}$
 - 2: $\mathcal{S}_i(t) := \text{smallest enclosing sphere of } \mathcal{R}_i(t)$
 - 3: $c_i(t) := \text{center of } \mathcal{S}_i(t)$
 - 4: Move towards $c_i(t)$ with speed 1, or stay on $c_i(t)$ if r_i is already positioned on it.
-

To show that CONT-3D-GTC is contracting, it must first be verified that connectivity of the visibility graph $UBG_t = (\mathcal{R}, E_t)$ is maintained at all times. The same reasoning that was used in the two-dimensional case by Li et al. [10] can also be applied here:

Lemma 12. *Let \mathcal{R} be a set of robots in the three-dimensional Euclidean space that follows the CONT-3D-GTC strategy. If $\{r_i, r_j\}$ is an edge in UBG_t at time t , then $\{r_i, r_j\}$ is an edge in $UBG_{t'}$ at $t' \geq t$. Thus, CONT-3D-GTC maintains the connectivity of UBG_t .*

Proof. Consider a robot r_i with neighborhood $\mathcal{R}_i(t)$ at time t . Let $Q_i(t)$ be the intersection of the unit balls of all robots in $\mathcal{R}_i(t)$. Since the SES of $\mathcal{R}_i(t)$ can have a radius of at most 1 and contains all robots in $\mathcal{R}_i(t)$, its center $c_i(t)$ must lie in $Q_i(t)$.

Consider some neighbor $r_j \in \mathcal{R}_i(t)$ of r_i and assume that there is some future point in time $t' > t$, such that $d(p_i(t'), p_j(t')) > 1$, i.e. r_i and r_j are no longer neighbors. Since the movement of robots is continuous, there must be some time $t^* \in [t, t']$, for which $d(p_i(t^*), p_j(t^*)) = 1$.

Now let L denote the intersection of the unit balls of r_i and r_j at time t^* . Any point in L is within distance at most 1 of both r_i and r_j . Furthermore L is a superset of both $Q_i(t^*)$ and $Q_j(t^*)$, meaning the target points $c_i(t^*)$ and $c_j(t^*)$ of both r_i and r_j also lie in L . Therefore, r_i and r_j can only move in the direction of points that are in distance at most 1 from both of them, meaning their distance can never exceed 1, creating a contradiction to the assumption that their distance is greater than 1 at time t' . \square \square

It remains to show that CONT-3D-GTC is a contracting strategy. This follows directly from Lemma 12 and Proposition 1, which states that the center of a SES is a convex combination of the points it encloses, meaning any target point computed by the strategy lies within the convex hull of the current configuration.

Theorem 13. *CONT-3D-GTC is a contracting, local strategy and thus gathers the robots in time $\mathcal{O}(\Delta \cdot n^{3/2})$.*

3.4 Tangential-Normal Strategies

Previously, we showed a runtime bound for a relatively general class of (not necessarily local) gathering strategies and introduced a concrete example in CONT-3D-GTC. However, when designing a local strategy, additional care has to be taken to maintain the visibility graph UBG_t to successfully solve the GATHERING problem. It would be useful to also have a relatively simple design criterion that ensures this property. For this purpose, we will focus on robots' *local convex hulls* and introduce the notion of tangential-normal strategies. Let $CH(\mathcal{R}_i(t))$ denote the local convex hull of robot r_i , i.e. the convex hull of r_i 's neighborhood. Furthermore, let $Adj_t(i)$ denote the set of robots that are adjacent to r_i on $CH(\mathcal{R}_i(t))$ if r_i lies on $CH(\mathcal{R}_i(t))$ itself. The main idea is to identify those velocity vectors that lead to a decrease in distance to all neighboring robots. These vectors are the normal vectors of *tangential planes*:

Definition 3. *Given a convex polyhedron $P \subset \mathbb{R}^3$ and a vertex $p \in P$. A tangential plane h_p w.r.t. P through p is a plane that only intersects P at the vertex p .*

Note that as long as P is actually convex, such a plane always exists and can – for example – be obtained by taking the plane through one of the faces adjacent to p and slightly rotating it. Based on this notion, we define the class of *tangential-normal* strategies in which the corner robots of local convex hulls move along the normal vectors of tangential planes:

Definition 4. *In the continuous time model, a gathering strategy for n robots is called tangential-normal if for every time t in which the robots have not yet gathered, each robot $r_i \in \mathcal{R}$ that is on a corner of its own local convex hull $CH(\mathcal{R}_i(t))$ moves with speed 1 along the normal vector of a tangential plane w.r.t. $CH(\mathcal{R}_i(t))$ through p_i while other robots do not move.*

The following lemma characterizes the normal vectors of tangential planes and will be used to show the desired properties of tangential-normal strategies.

Lemma 14. *Let p_i be a corner of a convex polyhedron P and let E_i be the set of edges of P adjacent to p_i . Then a plane h through p_i with normal vector \vec{n} is a tangential plane w.r.t. P if and only if for each edge $e \in E_i$, $\angle(\vec{n}, e) < \frac{\pi}{2}$*

Proof. First, note that by the convexity of P and since h only intersects with it in p_i , the entire rest of P lies on one side of h . However, if there was an edge e with $\angle(\vec{n}, e) \geq \frac{\pi}{2}$, this would mean that e lies on the opposite side of or directly on h , both of which are contradictions to h being a tangential plane.

For the other direction of the statement, let \vec{n} be a vector such that for each edge $e \in E_i$, $\angle(\vec{n}, e) < \frac{\pi}{2}$. This property now immediately yields that all edges $e \in E_i$ lie on the same side of the plane $h : \vec{n} \cdot (x - p_i) = 0$ defined by \vec{n} and the point p_i , making h a tangential plane w.r.t. P through p_i . \square

Theorem 15. *Let \mathcal{R} be a set of robots controlled by a tangential-normal strategy. Then, for each pair of robots $r_i, r_j \in \mathcal{R}$ and time t such that $\{r_i, r_j\}$ is an edge in UBG_t , $\{r_i, r_j\}$ is an edge in $UBG_{t'}$ for all $t' \geq t$. Thus, tangential-normal strategies maintain the connectivity of UBG_t .*

Proof. Let \mathcal{R} be a set of robots that follows a tangential-normal strategy. Consider a time t and a robot $r_i \in \mathcal{R}$ that lies on the corner of its own local convex hull $CH(\mathcal{R}_i(t))$ and let $v_i(t)$ be the normal vector of a tangential plane h_{p_i} w.r.t. $CH(\mathcal{R}_i(t))$. By Lemma 14, for each adjacent robot $r_j \in Adj_t(i)$, it holds that $\beta_{i,j}(t) < \frac{\pi}{2}$. Since the cosine is positive on the interval $[0, \frac{\pi}{2}]$, Lemma 5 yields that r_i contributes a strict decrease in distance to all of its neighbors. A neighbor r_j now

either does not move or also contributes a decrease in distance to r_i . This also means that for each pair of robots r_i and r_j that can see each other, the distance between r_i and r_j cannot increase, guaranteeing that the visibility graph remains connected. \square \square

Note however that while the tangential-normal property is a sufficient condition for ensuring connectivity, it is not a necessary condition. In particular, CONT-3D-GTC is not tangential-normal but is still able to maintain the connectivity of the visibility graph.

3.5 Move-on-Angle-Minimizer

Next, we introduce a strategy based on the tangential-normal criterion. It is based around the idea to find a movement vector that somehow causes a large decrease in distance to all neighbors. Since a smaller angle causes a greater decrease in distance (according to Lemma 5), one intuitive approach might be to find a movement vector that minimizes the maximal angle to all neighbors on the convex hull. To this end, the notion of an *angle minimizer* will be introduced. Let $V = \{\vec{v}_1, \vec{v}_2, \dots, \vec{v}_k\} \subset \mathbb{R}^3$ be a set of vectors that lie on one side of a plane through the origin. Then the vector $\vec{x}^* = \operatorname{argmin}_{\vec{x} \in \mathbb{R}^3} \max_{\vec{v}_i \in V} \angle(\vec{x}, \vec{v}_i)$ is called an angle minimizer of V .

Now, we define a strategy in which each robot that is a corner of its local convex hull r_i moves along the angle minimizer of the edges between itself and the robots in $\operatorname{Adj}_t(i)$. This strategy will be called MOVE-ON-ANGLE-MINIMIZER and is summarized in Algorithm 3.

Algorithm 3 MOVE-ON-ANGLE-MINIMIZER

- 1: $\mathcal{R}_i(t) := \{\text{positions of robots visible from } r_i, \text{ including } r_i \text{ at time } t\}$
 - 2: $CH(\mathcal{R}_i(t)) := \text{Convex hull of } r_i\text{'s neighborhood}$
 - 3: **if** r_i is on a corner of $CH(\mathcal{R}_i(t))$ **then**
 - 4: $\vec{x}^* = \operatorname{argmin}_{\vec{x} \in \mathbb{R}^3} \max_{r_j \in \operatorname{Adj}_t(i)} \angle(\vec{x}, p_j(t) - p_i(t))$
 - 5: r_i moves along \vec{x}^* with speed 1
 - 6: **else**
 - 7: r_i does not move
-

Note that if a robot r_i 's local convex hull is two-dimensional, the angle minimizer is identical to the angle bisector of the inner angle at r_i . Therefore, this strategy can also be viewed as a generalization of MOVE-ON-BISECTOR for two-dimensional continuous gathering [7], for which Kempkes et al. [3] could show an optimal gathering time of $\Theta(n)$.

It will now be shown that the presented strategy is both a tangential-normal and a contracting strategy. By Lemma 14 and the existence of a tangential plane, we already know that there is a possible movement vector that has an angle of less than $\pi/2$ to all neighbors on the local convex hull. Therefore, the same must hold for \vec{x}^* , immediately showing that MOVE-ON-ANGLE-MINIMIZER is a tangential-normal strategy.

Lemma 16. *MOVE-ON-ANGLE-MINIMIZER is a tangential-normal strategy.*

Computing \vec{x}^* In order to see that MOVE-ON-ANGLE-MINIMIZER is also a contracting strategy, we look at a method to compute the angle minimizer.

Let $\hat{v} = \vec{v}/\|\vec{v}\|$ denote the respective normalized vector of \vec{v} and let $\hat{V} = \{\hat{v} \mid \vec{v} \in V\}$ for a set V of vectors. Then the following holds:

Lemma 17. *Let $V \subset \mathbb{R}^3$ be a set of vectors that all lie on one side of a plane through the origin. The center \vec{c} of the smallest enclosing sphere of \hat{V} is an angle minimizer of V .*

Proof. Let \vec{x} be a vector such that $\angle(\vec{x}, \vec{v}_i) \leq \pi/2$ for all $\vec{v}_i \in V$. Such a vector exists, since there is a plane such that all \vec{v}_i lie on one side of this plane.

Now consider the normalized vectors $\hat{\vec{v}}_i$. They lie on the surface of the unit sphere centred on the origin. Let $C_{\vec{x}}$ be the minimal spherical cap centred on the vector \vec{x} such that all the $\hat{\vec{v}}_i$ lie on its surface. The vector $\hat{\vec{v}}_j \in \hat{V}$ with the maximal angle to \vec{x} lies on the edge of the base of $C_{\vec{x}}$. The maximal angle can now be computed using the radius r of $C_{\vec{x}}$ as $\angle(\vec{x}, \hat{\vec{v}}_j) = \sin^{-1} r$.

Since \sin^{-1} is monotonically increasing on the interval $[0, 1]$, finding the angle minimizer \vec{x}^* now amounts to finding the center \vec{c} of a spherical cap with minimal radius, which can be achieved by computing the smallest enclosing sphere of \hat{V} . \square \square

By applying the fact that the SES of \hat{V} is a convex combination of \hat{V} (Proposition 1), this lemma together with Lemma 16 immediately yields that MOVE-ON-ANGLE-MINIMIZER is also a contracting strategy.

Theorem 18. *MOVE-ON-ANGLE-MINIMIZER is a tangential-normal and a contracting strategy. Thus, it gathers the robots in time $\mathcal{O}(\Delta \cdot n^{3/2})$.*

References

- [1] Ando, H., Suzuki, Y., Yamashita, M.: Formation and agreement problems for synchronous mobile robots with limited visibility. In: Proceedings of the 1995 IEEE International Symposium on Intelligent Control, ISIC 1995. pp. 453–460. IEEE (Aug 1995). <https://doi.org/10.1109/ISIC.1995.525098>
- [2] Bhagat, S., Chaudhuri, S.G., Mukhopadhyaya, K.: Gathering of opaque robots in 3d space. In: Proceedings of the 19th International Conference on Distributed Computing and Networking, ICDCN 2018, Varanasi, India, January 4–7, 2018. pp. 2:1–2:10 (2018). <https://doi.org/10.1145/3154273.3154322>
- [3] Degener, B., Kempkes, B., Kling, 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>
- [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, June 4–6, 2011 (Co-located with FCRC 2011). pp. 139–148. ACM (2011). <https://doi.org/10.1145/1989493.1989515>
- [5] Elzinga, D.J., Hearn, D.W.: The minimum covering sphere problem. Manage. Sci. **19**(1), 96104 (Sep 1972). <https://doi.org/10.1287/mnsc.19.1.96>
- [6] Flocchini, P., Prencipe, G., Santoro, N. (eds.): Distributed Computing by Mobile Entities, Current Research in Moving and Computing, Lecture Notes in Computer Science, vol. 11340. Springer (2019). <https://doi.org/10.1007/978-3-030-11072-7>
- [7] 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.) Ant Colony Optimization and Swarm Intelligence. pp. 142–153. Springer Berlin Heidelberg, Berlin, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28646-2_13

- [8] Kling, P., Meyer auf der Heide, F.: Continuous protocols for swarm robotics. In: Distributed Computing by Mobile Entities, Current Research in Moving and Computing, pp. 317–334 (2019). https://doi.org/10.1007/978-3-030-11072-7_13
- [9] Li, S., Markarian, C., Meyer auf der Heide, F., Podlipyan, P.: A continuous strategy for collisionless gathering. In: Algorithms for Sensor Systems - 13th International Symposium on Algorithms and Experiments for Wireless Sensor Networks, ALGOSENSORS 2017, Vienna, Austria, September 7-8, 2017, Revised Selected Papers. pp. 182–197 (2017). https://doi.org/10.1007/978-3-319-72751-6_14
- [10] 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: Algorithms for Sensor Systems - 12th International Symposium on Algorithms and Experiments for Wireless Sensor Networks, ALGOSENSORS 2016, Aarhus, Denmark, August 25-26, 2016, Revised Selected Papers. pp. 62–79 (2016). https://doi.org/10.1007/978-3-319-53058-1_5
- [11] Poudel, P., Sharma, G.: Universally optimal gathering under limited visibility. In: Stabilization, Safety, and Security of Distributed Systems - 19th International Symposium, SSS 2017, Boston, MA, USA, November 5-8, 2017, Proceedings. pp. 323–340 (2017). https://doi.org/10.1007/978-3-319-69084-1_23
- [12] Tomita, Y., Yamauchi, Y., Kijima, S., Yamashita, M.: Plane formation by synchronous mobile robots without chirality. In: 21st International Conference on Principles of Distributed Systems, OPODIS 2017, Lisbon, Portugal, December 18-20, 2017. pp. 13:1–13:17 (2017). <https://doi.org/10.4230/LIPIcs.OPODIS.2017.13>
- [13] Yamauchi, Y., Uehara, T., Yamashita, M.: Brief announcement: Pattern formation problem for synchronous mobile robots in the three dimensional euclidean space. In: Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016. pp. 447–449 (2016). <https://doi.org/10.1145/2933057.2933063>

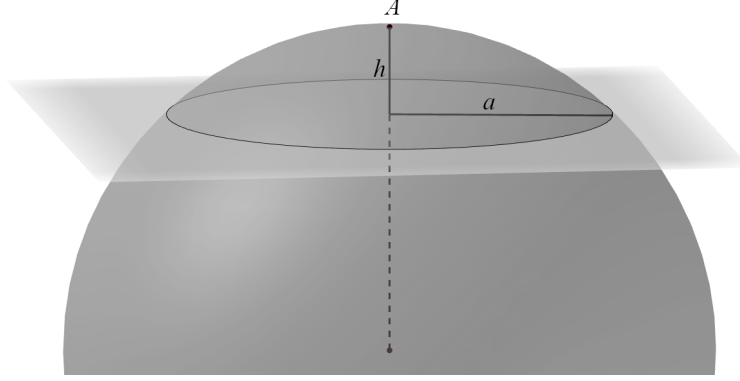


Figure 3: A plane dividing a sphere into two spherical caps. The smaller cap on top has apex A , height h and base radius a .

A Analysis of 3d-Go-To-The-Center

In this section, we show that the upper bound of $\mathcal{O}(n^2)$ that Degener et al. [4] presented for GO-TO-THE-CENTER also holds for 3D-GO-TO-THE-CENTER. As such, the analysis largely consists of a generalization of the proof by Degener et al. for three dimensions and therefore, each lemma provides an analogous result to a similar lemma used in the original proof for two dimensions.

Before beginning with the main analysis, some useful general properties of SESs will be stated. Two concepts that will be used throughout the proof are the notions of *convex combinations* and *spherical caps*.

Given a finite set of points $\mathcal{P} = \{p_1, \dots, p_n\}$, a point x is a convex combination of \mathcal{P} , if there are scalar coefficients $\alpha_i \geq 0$, such that $x = \sum_{i=1}^n \alpha_i p_i$ and $\sum_{i=1}^n \alpha_i = 1$. The set of all convex combinations of \mathcal{P} is identical to the convex hull of \mathcal{P} . A *spherical cap* is a region of a sphere cut off by a plane and is illustrated in Fig. A. It consists of a circular base with a radius a which is formed by the intersection of the plane and the sphere and a curved surface with an apex A . The distance between the base and the apex is called the height h of the spherical cap.

Proposition 1. [5] *Let \mathcal{S} be the smallest enclosing d -sphere (SES) of a point set $\mathcal{P} \subset \mathbb{R}^d$. Then the center c of \mathcal{S} is a convex combination of at most $d + 1$ points in \mathcal{P} that lie on the surface of \mathcal{S} . Especially,*

1. c lies in \mathcal{S}
2. c minimizes the maximum distance to the points in \mathcal{P} .

Definition 5. *Let \mathcal{S} be a sphere and \mathcal{P} a point set. A spherical cap of \mathcal{S} whose curved surface does not contain any points from \mathcal{P} is called a point-free cap.*

Using this definition, the following Lemma can be stated about SESs:

Lemma 19. *Let \mathcal{S} be the SES with radius r of a set of $n \geq 2$ points. Then there is no point-free cap of \mathcal{S} with height h , such that $h > r$.*

Proof. Let \mathcal{S} be the SES of a set of $n \geq 2$ points with center c and radius r . Assume there is a point-free cap with height $h > r$.

Without loss of generality, rotate and translate the coordinate system such that c lies at the origin and the apex of the point-free cap lies on the positive x -axis. Let c_x be the x -component of c . Then for any point p on the surface of \mathcal{S} , its x -component p_x is at most

$$p_x < c_x - (h - r) < c_x. \quad (6)$$

Therefore, c cannot be a convex combination of points on the surface of \mathcal{S} . By Prop. 1, c cannot have been the center of a SES, leading to a contradiction. \square \square

A.0.1 Progress Measures

In order to determine the progress of the gathering, two measures will be used. Firstly, consider what happens if two robots move to the exact same position at the same time: Since the strategy is deterministic, the two robots will always observe the same neighborhood and therefore always compute the same target point. We say two such robots have *merged*. Starting with n robots, there can clearly be at most $n - 1$ rounds with merging events until all robots have gathered. This number of merges will be used as the first progress measure.

Now let \mathcal{N}_t be the global SES around a center \mathbf{M}_t with radius R_t in some round $t \geq 0$. The radius R_t can be used as a second progress measure: The robots always move towards the centers of their respective (local) SES. By Prop. 1, these centers always lie within the convex hulls of the respective neighborhoods and therefore also within the global convex hull. Since robots do not leave the global convex hull, they also cannot leave \mathcal{N}_t , meaning R_t can not increase with time.

On the other hand, once R_t is smaller than $1/2$, all robots are within distance of at most 1 of each other, meaning each robot sees every other robot. At this point, they will all compute the same target point \mathbf{M}_t and be able to move to it since the limiting spheres $\mathcal{B}_j(t)$ with radii $1/2$ must contain \mathbf{M}_t for all pairs of robots r_i and r_j .

The overall idea of the proof is now to show that during a constant number of rounds, either two robots merge, or R_t decreases by at least $\Omega(1/n)$. Since the initial radius R_0 is $\mathcal{O}(n)$ if the initial configuration is connected, this yields a total gathering time of $\mathcal{O}(n^2)$ rounds.

During the rest of the analysis, consider some arbitrary but fixed round t_0 and let \mathcal{N} and R denote \mathcal{N}_{t_0} and R_{t_0} respectively. In order to show the progress of the two measures mentioned above, the analysis will focus on a certain region of \mathcal{N} . Let P be an arbitrary point on the surface of \mathcal{N} and define a spherical cap C with apex P as follows (see Fig. A.0.1 for an illustration of the construction): Choose the height h of C , such that the slant height of the inscribed cone of C is $1/8$, i.e. such that $h^2 + a^2 = (1/8)^2$. Note that this causes the radius a to be bounded by $1/8$. Therefore the diameter of C as a whole is at most $1/4$. C is then separated into two parts by intersecting \mathcal{N} with another plane that is parallel to the base of C , creating another spherical cap C_1 with height $h/2$ that also has P as its apex. The remaining spherical segment $C \setminus C_1$ is called C_2 .

The goal is to now show that within two consecutive rounds t_0 and $t_0 + 1$, either at least two robots merge, or all robots leave C_1 . In order to quantify the progress made in case the latter happens, a bound on h is needed, which will be provided by the following lemma:

Lemma 20. *The spherical cap C has a height h of at least $\frac{1}{64\pi \cdot R} \in \Omega(\frac{1}{n})$.*

Proof. The proof proceeds very similarly to that of the analogous lemma by Degener et al. [4] for the two-dimensional case. The goal is to use the angle α between the height h of C and the (known) slant height $(1/8)$ of the inscribed cone of C (see Fig. A.0.1) to compute a bound on h . In order to do this, consider a circle K with radius R , with the same midpoint as \mathcal{N} and containing P . The angle α can be considered to be half of the internal angle of a regular convex polygon with side length $1/8$ and whose m vertices lie on K . Since the circumference of K is $2\pi R$, the number

of vertices m is bounded by $16\pi R$. The sum of angles of such a convex polygon is now given by $(m-2)\pi$. This yields the following bound on α :

$$\begin{aligned} 2\alpha &\leq \frac{(16\pi R - 2)\pi}{16\pi R} = \pi - \frac{1}{8R} \\ \Leftrightarrow \alpha &\leq \frac{\pi}{2} - \frac{1}{16R} \end{aligned}$$

Using the fact that $\cos(x) \geq -\frac{2}{\pi}x + 1$ for $0 \leq x \leq \pi/2$, we can now obtain a bound on h :

$$\begin{aligned} h &= \frac{\cos(\alpha)}{8} \geq \frac{1}{8} \cdot \left(-\frac{2}{\pi}\alpha + 1 \right) \\ &\geq \frac{1}{8} \cdot \left(-\frac{2}{\pi} \cdot \left(\frac{\pi}{2} - \frac{1}{16R} \right) + 1 \right) \\ &= \frac{1}{64\pi R} \end{aligned}$$

Since \mathcal{N} is a SES, its radius R is at most $n/2$. Therefore, we have $h \in \Omega(\frac{1}{n})$. □ □

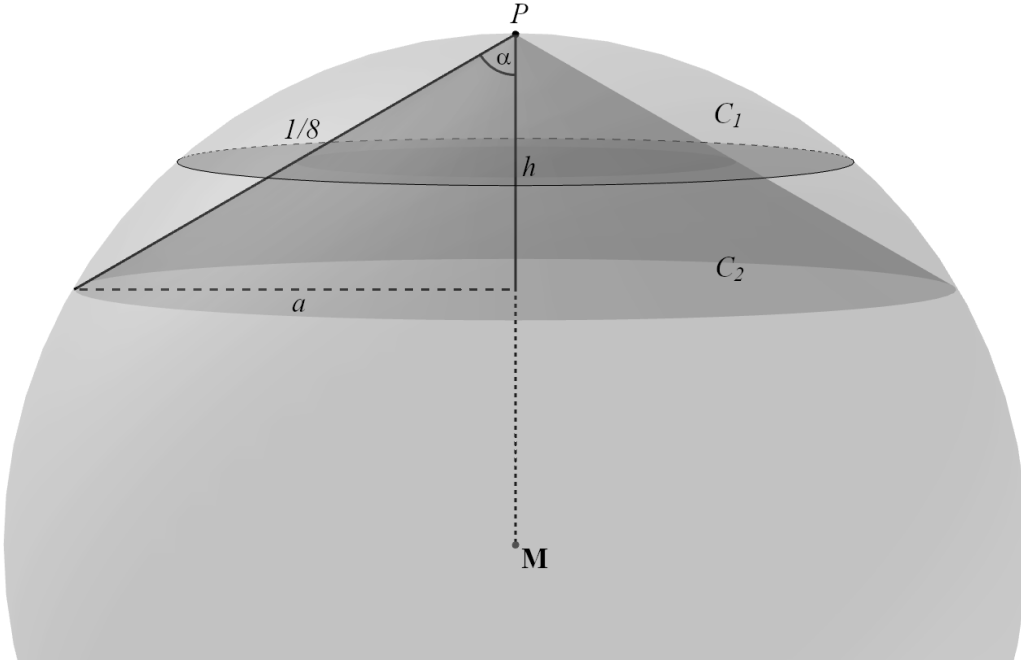


Figure 4: Illustration of the spherical cap C with apex P , consisting of the smaller cap C_1 and the spherical segment C_2 . The height h is chosen such that the inscribed cone of C has a slant height of $1/8$.

A.0.2 Properties of Target Points

First, some properties of SES that can be computed by robots in and around C will be given. These will be useful later on to identify situations in which robots leave C , as well as ensuring that no further robots enter it.

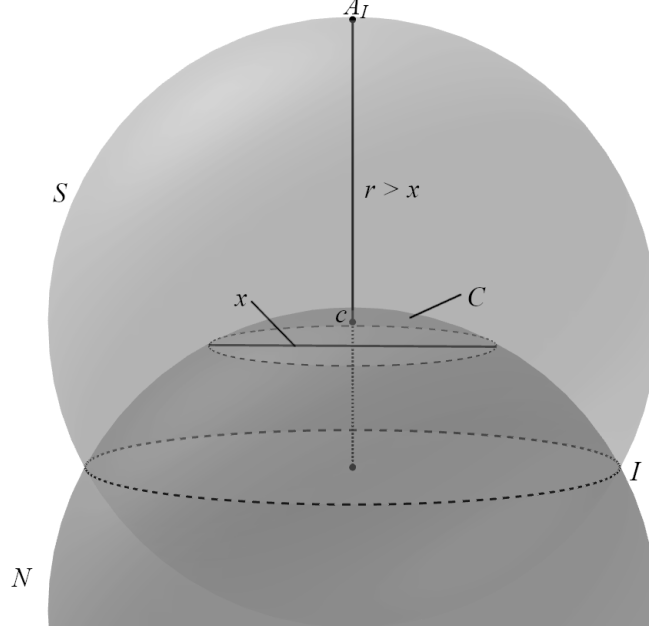


Figure 5: Illustration of the setting of Lemma 21. A sphere with its center c in the spherical cap C and a radius r that is greater than the diameter of x of C cannot be a SES of points in \mathcal{N} .

Lemma 21. *Let x be the diameter of a spherical cap C of \mathcal{N} . Then any sphere S with its center c in C and radius $r > x$ cannot be a SES of points in \mathcal{N} .*

Proof. The setting of the lemma is described in Fig. A.0.2. Since the center of S lies in C and the radius r of S is greater than the diameter x of C , every intersection point of S and \mathcal{N} lies outside of C . The (circular) intersection of S and \mathcal{N} , $I = S \cap \mathcal{N}$, therefore lies completely outside of S . Now let C_I be the spherical cap of S whose curved surface lies outside of \mathcal{N} and that has I as its base. Let h_I be its height and A_I be its apex. Note that C_I is therefore a point-free cap w.r.t. any set of points lying in \mathcal{N} .

Now consider the height h_I of C_I : Since I lies outside of C , so does its center c_I . The line segment between A_I and c_I therefore has to pass through c , which lies within C . Thus, we have that $h_I > r$. By Lemma 19, this means that S cannot have been a SES of points in \mathcal{N} . $\square \quad \square$

Since the diameter of C is at most $1/4$, this immediately yields the following corollary:

Corollary 22. *The radius of a SES of a point set $\mathcal{P} \subseteq \mathcal{N}$ with its center in C is at most $1/4$.*

The next lemma will be used to show that a robot with a neighbor that is positioned far away from C cannot compute a target point in C .

Lemma 23. *Let $\mathcal{P} \subseteq \mathcal{N}$ be a set of points and let \mathbf{A} be a point in C (not necessarily in \mathcal{P}). If there is a point $\mathbf{B} \in \mathcal{P}$ in distance more than $1/2$ from \mathbf{A} , then the center of the SES of \mathcal{P} cannot lie in the cap C .*

Proof. Assume the SES S of \mathcal{P} has its center c in C . Corollary 22 shows that the radius of S can be at most $1/4$ and since \mathbf{B} lies in S , we have $d(c, \mathbf{B}) \leq 1/4$. Since \mathbf{A} and c both lie in C (which has a diameter of at most $1/4$), we also have $d(c, \mathbf{A}) \leq 1/4$. Applying the triangle inequality yields:

$$d(\mathbf{A}, \mathbf{B}) \leq d(c, \mathbf{A}) + d(c, \mathbf{B}) \leq \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$$

Which is a contradiction to $d(\mathbf{A}, \mathbf{B}) > 1/2$, meaning c cannot lie in C . □ □

The next lemma concerns local configurations of robots consisting of only a single robot in C and a number of robots outside of C . It shows that in such a scenario, the target point computed by the strategy must lie outside of C_1 .

Lemma 24. *The center of the SES \mathcal{S} of a non-empty point set $\mathcal{P} \subseteq \mathcal{N} \setminus C$ and a point $\mathbf{A} \in C$ cannot lie in the spherical cap C_1 .*

Proof. By Prop. 1, the center c of \mathcal{S} is a convex combination of at most 4 points lying on the surface of \mathcal{S} . In other words, c is the center of the circumscribed circle of either a line, a triangle or a tetrahedron. There are multiple cases that can occur regarding those points defining \mathcal{S} .

First, consider the case that there is no subset of points $\mathcal{P}' \subseteq \mathcal{P} \cup \{\mathbf{A}\}$ on the surface of \mathcal{S} containing \mathbf{A} , such that c is a convex combination of points from \mathcal{P}' . Then c must lie outside of C , since it is a convex combination of points outside of C .

Now consider the case that \mathbf{A} is one of the points defining c , i.e. it is part of a set of points \mathcal{P}' such that c is a convex combination of \mathcal{P}' . If \mathbf{A} does not lie in C_1 , then c once again cannot lie in C_1 either. Now assume that $\mathbf{A} \in C_1$ and consider multiple sub-cases depending on the cardinality of \mathcal{P}' : The first two cases are completely analogous to the proof of the two-dimensional version of the Lemma shown by Degener et al. [4].

- Case $|\mathcal{P}'| = 2$: \mathcal{S} is defined by \mathbf{A} and another point $P \in \mathcal{P}$. Since \mathcal{S} is a SES, c is the midpoint between \mathbf{A} and P . Between \mathbf{A} and P lies the spherical segment C_2 , with the same height as C_1 , in which \mathbf{A} lies. Therefore, the midpoint between them cannot lie in C_1 .
- Case $|\mathcal{P}'| = 3$: \mathcal{S} is the circumscribed circle of a triangle $\Delta \mathbf{A}P_1P_2$ formed by \mathbf{A} and two other points P_1 and P_2 . The center of \mathcal{S} can now be computed by determining the intersection of the perpendicular bisectors of two of the triangle's sides, say $\overline{\mathbf{A}P_1}$ and $\overline{\mathbf{A}P_2}$. Since \mathbf{A} lies inside C_1 and the other two points lie outside of C , the centers of those two sides cannot lie in C_1 . Furthermore, c must lie inside of the triangle as it is a convex combination of its corners, meaning it is acute. Therefore, the intersection point of the perpendicular bisectors cannot lie in C_1 .
- Case $|\mathcal{P}'| = 4$: \mathcal{S} is the circumscribed sphere of a tetrahedron T formed by the points \mathbf{A} , P_1 , P_2 and P_3 . Similarly to the previous case, c lies at the intersection of the perpendicular bisector planes of three edges of T . Choose those three edges as those adjacent to \mathbf{A} . Their midpoints once again must lie outside of C_1 , since the points P_1 , P_2 and P_3 lie outside of C . Since c has to lie inside of T , and the three edges are all adjacent to the same point \mathbf{A} , their bisector planes can only meet below their midpoints, outside of C_1 .

□

□

A.0.3 Hindering Robots

In the previous subsection, several properties concerning the target points of robots have been established. However, due to the second part of the strategy, robots are not always able to reach their target points. In particular, if a robot r_i has a neighbor r_j such that r_i 's target point lies outside of the *limit sphere* $\mathcal{B}_j(t)$ of r_j , then r_i will be stopped by that limit sphere. If this is the case, we say that r_j *hinders* r_i from reaching its target point.

Next, it will be established that robots always reach their target points if they lie within C .

Lemma 25. *Robots that compute a target point in C cannot be hindered from reaching it by any other robot.*

Proof. Let r_i be a robot that computes a target point c within C , which is the center of the SES \mathcal{S}_i . Assume there is a robot r_j in the neighborhood of r_i that hinders r_i from reaching c .

By Corollary 22, we know that \mathcal{S}_i has a radius of at most $1/4$. Since r_j 's position must lie within \mathcal{S}_i , it can have a distance of at most $1/2$ to r_i . Therefore, the midpoint m_j between r_i and r_j 's positions is within a distance of at most $1/4$ from r_i . This leaves a distance of at least $1/2 - 1/4 = 1/4$ that r_i can move freely without leaving the limit sphere of r_j . As \mathcal{S}_i has a radius of at most $1/4$, this is enough to reach c , meaning r_j cannot have hindered r_i . \square \square

Similarly, robots cannot be hindered from leaving C , as is shown by the following lemma.

Lemma 26. *Robots cannot be hindered from leaving C by any other robot.*

Proof. Let r_i be a robot in C that computes a target point c outside of C and let m be the point where r_i would leave C . Note that $d(p_i, m) \leq 1/4$, since the diameter of C is at most $1/4$.

Now assume that there is a robot r_j hindering r_i from leaving C . First of all, r_j must be a neighbor of r_i , meaning it is within distance 1 of r_i as well as c . Now, let $m_j = \frac{p_i + p_j}{2}$ be the midpoint between r_i and r_j . For r_j to hinder r_i from reaching m , m must lie outside of r_j 's limit sphere, i.e. $d(m, m_j) > 1/2$. Furthermore, let $p'_i = 2m - p_i$ be the reflection of p_i about m . Note that $m = \frac{p'_i + p_i}{2}$. Putting these together yields:

$$\begin{aligned} d(m, m_j) &= \|m - m_j\| &> \frac{1}{2} \\ \iff \left\| \frac{p'_i}{2} + \frac{p_i}{2} - \frac{p_i}{2} - \frac{p_j}{2} \right\| &> \frac{1}{2} \\ \iff \|p'_i - p_j\| &> 1 \\ \iff d(p_j, p'_i) &> 1 \end{aligned}$$

To summarize, we now have the following three constraints on the position of p_j :

$$d(p_j, p_i) \leq 1 \tag{7}$$

$$d(p_j, c) \leq 1 \tag{8}$$

$$d(p_j, p'_i) > 1 \tag{9}$$

From this we can conclude that c must lie in between p_i and p'_i as follows: Assume p'_i lies between c and p_i instead. By the constraints (7) and (8), p_j must lie within the overlap of two balls of radius 1 centered around p_i and c . However, it must lie *outside* of the ball of radius 1 centered around p'_i (9). Since p'_i lies on the line segment between p_i and c by assumption, the entire overlap of the first two balls must be contained within the last ball around p'_i , leaving no viable positions for p_j . Therefore, p'_i cannot lie between p_i and c and thus, c must lie between p_i and p'_i instead.

Based on what is now known about the position of c , another spherical segment can be defined that has to contain p'_i and c : Let C_3 be the spherical segment of \mathcal{N} below C (see Fig. A.0.3). It is defined by the base of C and another plane l which is parallel to and in distance h of the base of C . The diameter of C_3 is now the diameter of the intersection of l and \mathcal{N} and can be bounded as follows: Once more consider the inscribed cone of C . Now extend this cone until it intersects with l . Note that this extended cone now has height $2h$ and slant height $1/4$ (since the slant height of the initial inscribed cone was $1/8$ by construction). Therefore, the radius of its base is also at most

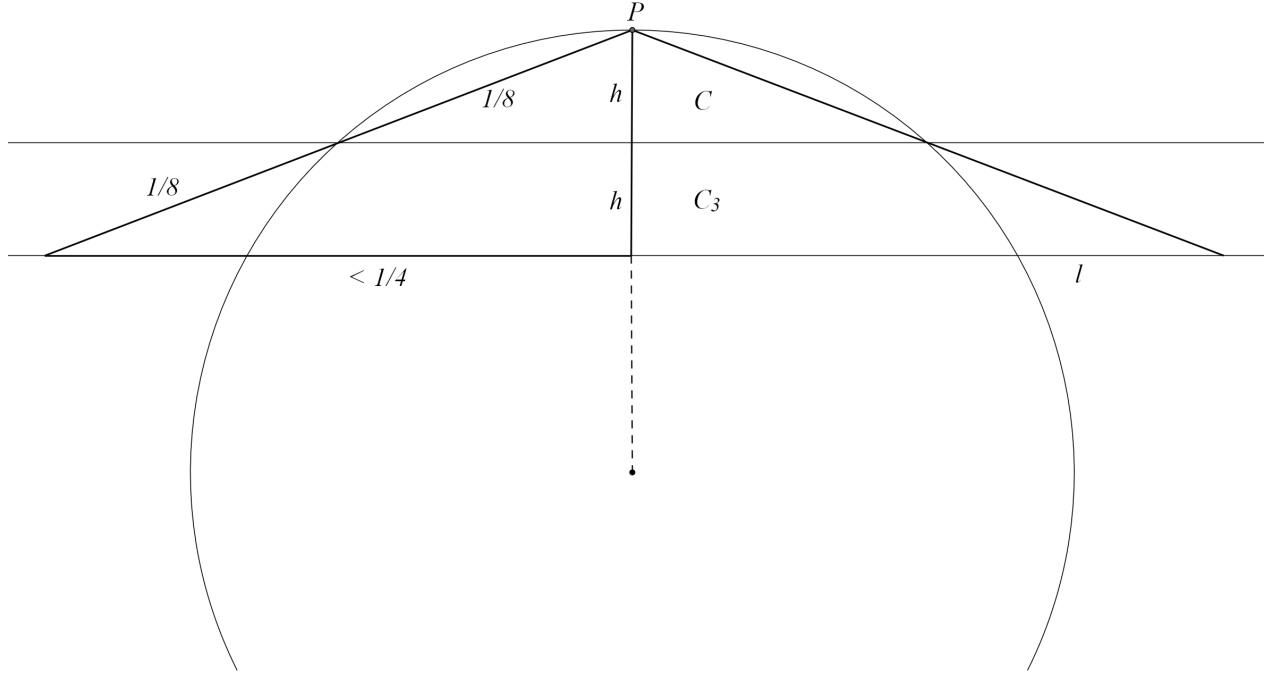


Figure 6: A cross-section of \mathcal{N} illustrating the construction of the spherical segment C_3 . It lies below and has the same height as C and is defined by the base of C and the plane l . Its diameter is at most $1/2$.

1/4. Since the cone intersects \mathcal{N} at the base of C , it must intersect l outside of \mathcal{N} , meaning the base of C_3 is completely contained within the base of the cone. Therefore, the diameter of C_3 can be at most $1/2$.

Note that since p'_i was defined as the reflection of p_i about a point on the base of C , and C_3 has the same height as C , p'_i and therefore also c must be contained in C_3 . Thus, we have $d(c, p'_i) \leq 1/2$. Using Lemma 21 and the diameter of C_3 , the radius of the SES corresponding to the target point c can now be further bounded to be at most $1/2$, i.e. $d(p_j, c) \leq 1/2$. Putting these two bounds together and applying the triangle inequality yields:

$$d(p_j, p'_i) \leq d(p_j, c) + d(c, p'_i) \leq \frac{1}{2} + \frac{1}{2} = 1$$

This is a contradiction to $d(p_j, p'_i) > 1$ (9), meaning r_j cannot have hindered r_i from leaving C . □

A.0.4 The Upper Bound

Now that the necessary preliminaries have been established, putting them together to acquire the main result works completely analogously to the proof for the two-dimensional case shown by Degener et al. [4].

Lemma 27. *Let \mathcal{P} be a set of robots in round t_0 that are all positioned in or compute a target point in C and that all have pairwise different neighborhoods. Then at most one of those robots is in C at the beginning of the next round.*

Proof. Let r_i be a robot whose set of neighbors $\mathcal{R}_i(t_0)$ is minimal, i.e. no other robot $r_j \in \mathcal{P} \setminus \{r_i\}$ has a neighborhood $\mathcal{R}_j(t_0)$ that is a strict subset of $\mathcal{R}_i(t_0)$. Therefore, all other robots r_j have a neighbor r_k that is not seen by r_i , i.e. $d(r_i, r_k) > 1$.

First consider the case that r_i is positioned in C : In this case Lemma 23 can be applied to show that no neighbor r_j of r_i can compute a target point in C by choosing $\mathcal{R}_j(t_0)$ as the point set \mathcal{P} , the position of r_i as the point \mathbf{A} and the position of r_k as the point \mathbf{B} in distance greater than $1/2$ of \mathbf{A} . Note that any robot outside of C that could potentially compute a target point in C must be within distance $1/4$ of C by Corollary 22 and is therefore at most a distance of $1/2$ away from r_i and thus its neighbor. It follows that only r_i could possibly remain in C .

On the other hand, if r_i is positioned outside of C , r_i computes a target point in C by the definition of \mathcal{P} . By Corollary 22, the SES of $\mathcal{R}_i(t_0)$ now has a radius of at most $1/4$, meaning r_i is also at most a distance of $1/4$ away from C and also in distance at most $1/2$ from any point in C . Now again consider a robot $r_j \in \mathcal{P} \setminus \{r_i\}$ with a neighbor r_k that is unseen by r_i . For this robot r_k , we must have $d(r_k, C) > 1/2$, otherwise it would be a neighbor of r_i , since it would be in distance 1 of it by applying the triangle inequality:

$$d(r_k, C) + d(r_i, C) \leq 1 < d(r_i, r_k)$$

From this, it can be concluded that r_j 's target point cannot lie in C : If it was in C , the radius of $\mathcal{R}_j(t_0)$'s SES would also be at most $1/4$ by Corollary 22, meaning it could not contain r_k . This means that r_j is positioned in C by definition of \mathcal{P} and computes a target point outside of C . By Lemma 26, r_j cannot be hindered from leaving C . Since this holds for all robots except for r_i , r_i is once again the only robot that might remain in C in the following round $t_0 + 1$. \square \square

Using this lemma, it is now finally possible to show that the progress measures mentioned earlier in this section always improve during a pair of consecutive rounds:

Lemma 28. *If $R_t \geq 1/2$, either there are robots that merge in round t or after rounds t and $t + 1$, the cap C_1 does not contain any robots.*

Proof. Consider all robots that are positioned in or compute a target point in C . We can distinguish two types of robots: Firstly, there are robots r_i that have a neighbor r_j with the same neighborhood, i.e. $\mathcal{R}_i(t) = \mathcal{R}_j(t)$. This means that r_i and r_j compute the same target point. If there is such a pair of robots that compute a common target point in C , then by Lemma 25, they both reach it and thus merge and fulfill the Lemma to be proven. On the other hand, if such a pair does not exist, then all robots sharing neighborhoods with other robots must compute target points outside of C and are not hindered from leaving C (Lemma 26) if they were positioned inside of it, meaning none of them can remain within C at the end of the round.

It remains to consider the set of robots that all have pairwise different neighborhoods and also either lie in or compute a target point in C . This is the exact situation described by Lemma 27, meaning at most one of these robots can remain within C after round t .

Therefore, in the beginning of round $t + 1$ and if no robots merged in round t , at most one robot r_i remains within C . If it lies in C_2 , we are done. Otherwise, if it lies in C_1 , only r_i itself and its neighbor could possibly compute a target point in C_1 . However, by Lemma 24, this cannot happen and by Lemmas 25 and 26, r_i cannot be hindered from leaving C_1 at which point no robots remain in C_1 . \square \square

Using this lemma now yields the main result.

Theorem 4. *Given n robots in a connected starting configuration $\mathcal{P} \in \mathbb{R}^3$ in the Euclidean space, 3D-GTC gathers the robots in $\mathcal{O}(n^2)$ rounds.*

Proof. Fix an arbitrary round $t_0 \geq 0$. Lemma 28 holds for any spherical cap C with an arbitrary point P on the boundary of the global SES \mathcal{N}_{t_0} as its apex. Therefore, either at least two robots merge or all robots are within distance $h/2$ of the boundary of \mathcal{N}_{t_0} at the beginning of round $t_0 + 2$. This means that the radius of \mathcal{N}_{t_0+2} is at least $h/2$ smaller than that of \mathcal{N}_{t_0} since it is a SES. By Lemma 20, we have:

$$\frac{h}{2} \geq \frac{1}{128\pi \cdot R_{t_0}} \geq \frac{1}{128\pi \cdot R_0} \geq \frac{1}{128\pi \cdot n}$$

Consequently, it takes at most $\lceil 2 \cdot 128 \cdot \pi \cdot n^2 \rceil$ rounds without merging robots until the radius is less than $1/2$, at which point all robots can see each other, compute the same target and move towards it in a single round.

Therefore, there can overall be either at most $n - 1$ rounds with merges or $\mathcal{O}(n^2)$ rounds without merges until all robots have gathered. □ □