Gathering in the Plane of Location-Aware Robots in the Presence of Spies

Jurek Czyzowicz[14], Ryan Killick[2], Evangelos Kranakis[25], Danny Krizanc[3], and Oscar Morale-Ponce[4]

[1] Départemant d'informatique, Université du Québec en Outaouais, Canada.
[2] School of Computer Science, Carleton University, Ottawa, Ontario, Canada.
[3] Department of Mathematics & Computer Science, Wesleyan University, Middletown CT, USA.
[4] Department of Computer Science, California State University, Long Beach, CA, USA.
[5] Research supported in part by NSERC Discovery grant.

Abstract. A set of mobile robots (represented as points) is distributed in the Cartesian plane. The collection contains an unknown subset of byzantine robots which are indistinguishable from the reliable ones. The reliable robots need to gather, i.e., arrive to a configuration in which at the same time, all of them occupy the same point on the plane. The robots are equipped with GPS devices and at the beginning of the gathering process they communicate the Cartesian coordinates of their respective positions to the central authority. On the basis of this information, without the knowledge of which robots are faulty, the central authority designs a trajectory for every robot. The central authority aims to provide the trajectories which result in the shortest possible gathering time of the healthy robots. The efficiency of a gathering strategy is measured by its competitive ratio, i.e., the maximal ratio between the time required for gathering achieved by the given trajectories and the optimal time required for gathering in the offline case, i.e., when the faulty robots are known to the central authority in advance. The role of the byzantine robots, controlled by the adversary, is to act so that the gathering is delayed and the resulting competitive ratio is maximized.

The objective of our paper is to propose efficient algorithms when the central authority is aware of an upper bound on the number of byzantine robots. We give optimal algorithms for collections of robots known to contain at most one faulty robot. When the proportion of byzantine robots is known to be less than one half or one third, we provide algorithms with small constant competitive ratios. We also propose algorithms with bounded competitive ratio in the case where the proportion of faulty robots is arbitrary.

1 Introduction

1.1 The background

A collection of mobile robots need to meet at some point of the geometric environment. This task, known as *gathering* or *rendezvous*, has been extensively investigated in the past. The gathering may be necessary, e.g., to coordinate a future task or to exchange previously acquired information.

In most formerly studied cases, robots have limited knowledge about the environment and they do not know the positions of the other robots. In the present paper, the robots are distributed in the two-dimensional Cartesian plane. They are equipped with GPS devices and they can wirelessly communicate their positions to the central authority. The central authority then informs each individual robot of the trajectory it is to follow in order to meet. However, the team of reliable robots has been contaminated with "spies" - a subset of byzantine robots, indistinguishable from the original ones, controlled by an omnipotent adversary. The role of the faulty robots is simple – delay the gathering of the reliable ones for as long as possible. A byzantine robot may report a wrong position, fail to report any, or fail to follow its assigned route. As the central authority does not recognize which robots are byzantine, it sends the travel instructions to all of them.

Our goal is to design a strategy resulting in gathering of all reliable robots within the smallest possible time. We attempt to minimize the *competitive ratio* – the ratio of the time required to achieve gathering of the reliable robots, to the time required for such gathering to occur under the assumption that the reliable robots were known in advance.

1.2 The model and the problem

A collection $\mathcal{S}$ of n mobile robots move at maximum unit speed within the two-dimensional plane. It is assumed that each robot in $\mathcal{S}$ is equipped with a GPS device so it is aware of a pair of Cartesian coordinates representing its current location in the plane.

We consider the problem of gathering an unknown subset $\mathcal{N} \subseteq \mathcal{S}$ of robots. The robots of $\mathcal{N}$ need to arrive at some time at a same point on the plane in order to complete some given task.We refer to this set $\mathcal{N}$ of at least $n - F$ robots as the set of reliable robots and define $\mathcal{F} = \mathcal{S} \setminus \mathcal{N}$ of $f \leq F$ robots as the set of byzantine robots. We call this problem of gathering all reliable robots from a collection containing at most F byzantine robots the $Gather(n, F)$ problem.

At the beginning, all robots in $\mathcal{S}$ send a single message recording their starting positions to the central authority. In turn, the central authority computes a set of trajectories instructing each robot how to time their respective movements in order to achieve gathering. At this point the robots follow the trajectories provided.

The movement continues until all reliable robots meet for the first time. We imagine a successful gathering as a meeting of robots possessing pieces of information allowing them to solve some puzzle. As long as all pieces are disassembled, the puzzle remains unsolved, and the identification of useful or invalid information is not possible.

The byzantine robots may report incorrect initial locations, which can potentially adversely affect the robots' trajectories. Clearly, this results in byzantine robots not being able to follow the assigned trajectories. However, as long as all reliable robots complete their trajectories, the schedule must lead to their gathering.

The trajectories designed by the central authority are computed uniquely on the basis of the reported set of robot positions and possibly using the knowledge of the upper bound on the number of byzantine robots. Once the robots start their movements, no adaptation to our algorithm is ever possible as no extra information may be obtained. We assume that the adversary knows in advance our algorithm and it will put the byzantine robots in the positions which result in the worst possible competitive ratio.

We note that the requirement of a central authority may be removed by allowing the robots to instead broadcast their initial positions to all other robots. In this situation all robots compute the same set of trajectories using the same algorithms.

We are interested in developing algorithms solving the $Gather(n, F)$ problem which are optimal in terms of the competitive ratio for a given initial configuration $\mathcal{S}$ of n robots, at most F of which are byzantine. We define the competitive ratio $\mathrm{CR}_{n,F}(\mathcal{A}, \mathcal{N})$ of an algorithm $\mathcal{A}$ for the specific subset $\mathcal{N}$ of the input $\mathcal{S}$ as the ratio of the time $T_{\mathcal{A}}(\mathcal{N})$ – the time of the first gathering of all robots belonging to $\mathcal{N}$ – divided by $T_*(\mathcal{N})$ – the minimal time necessary to gather the robots in $\mathcal{N}$, i.e. $\mathrm{CR}_{n,F}(\mathcal{A}, \mathcal{N}) = \frac{T_{\mathcal{A}}(\mathcal{N})}{T_*(\mathcal{N})}$. We also define the *overall* competitive ratio $\widehat{\mathrm{CR}}_{n,F}(\mathcal{A}, \mathcal{S})$ of an algorithm $\mathcal{A}$ with input $\mathcal{S}$ as the maximal $\mathrm{CR}_{n,F}$ over any subset $\mathcal{N}$ of $\mathcal{S}$, i.e. $\widehat{\mathrm{CR}}_{n,F}(\mathcal{A}, \mathcal{S}) = \max_{\mathcal{N} \subseteq \mathcal{S}} \mathrm{CR}_{n,F}(\mathcal{A}, \mathcal{N})$. We further define the *optimal* competitive ratio $\overline{\mathrm{CR}}_{n,F}(\mathcal{S})$ for an input $\mathcal{S}$ as the minimal $\widehat{\mathrm{CR}}_{n,F}(\mathcal{A}, \mathcal{S})$ for any algorithm $\mathcal{A}$, i.e. $\overline{\mathrm{CR}}_{n,F}(\mathcal{S}) = \min_{\mathcal{A}} \widehat{\mathrm{CR}}_{n,F}(\mathcal{A}, \mathcal{S})$. For ease of presentation we will often drop the subscripts n and F when they are implied by context.

We define an optimal algorithm $\overline{\mathcal{A}}$ solving the $Gather(n, F)$ problem as any algorithm satisfying

$$\mathrm{CR}_{n,F}(\overline{\mathcal{A}}, \mathcal{S}) = \overline{\mathrm{CR}}_{n,F}(\mathcal{S}), \ \forall \, \mathcal{S}. \tag{1}$$

1.3 Our results

We provide algorithms with constant competitive ratio for all but a small bounded region in the space of possible n and F pairs. In doing so we demonstrate that having knowledge of the upper bound of the number of byzantine robots in the subset (represented by the parameter F) permits fine-tuning of the gathering algorithm, resulting in better competitive ratios.

In Section 2 we consider the gathering problem for collections involving only a single byzantine robot. After developing insight into the problem we give a gathering algorithm that is optimal for any number of robots, at most one of which is byzantine. For the boundary case of three robots, one of which is byzantine, we give a closed form expression for the competitive ratio.

Section 3 presents two algorithms with small constant competitive ratio when the number of byzantine robots is bounded by a small fraction of n. Specifically, we give algorthms with competitive ratios of 2 and $2\sqrt{2}$ when $F < \lceil n/3 \rceil$ and $F < \lceil n/2 \rceil$ respectively.

Finally, in Section 4, we give two gathering algorithms solving the problem for any n and any F. The competitive ratio of one of these algorithms is constant, while the other is bounded by $F + 2$.

We summarize the results of the paper in Table 1 and Figure 1.

Table 1: Summary of competitive ratio bounds for various algorithms.

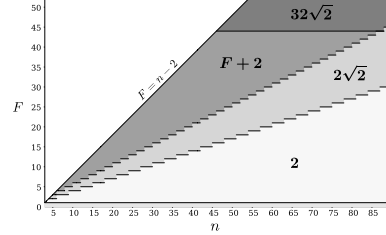| F	Upper-bound	Reference
1	optimal	Alg. 4
$\leq \lceil n/3 \rceil$	2	Alg. 5
$\leq \lceil n/2 \rceil$	$2\sqrt{2}$	Alg. 6
$> \lfloor 32\sqrt{2} \rfloor - 2$	$32\sqrt{2}$	Alg. 7
$\leq \lfloor 32\sqrt{2} \rfloor - 2$	$F + 2$	Alg. 8

Fig. 1: Competitive ratio bounds for various regions of the space of possible n and F pairs.

1.4 Related work

The gathering problem was originally introduced in [34] as a version of pattern formation (see also [21]). In operations research, Alpern [2,3] considers the gathering of two robots, referred to as the *rendezvous* problem (cf. [31]). Both problems are central in theoretical computer science. The rich related literature is due to the large variety of studied settings: deterministic and randomized, synchronous and asynchronous, for labeled and anonymous agents, in graphs and geometric environments, for same-speed or distinct maximal speed agents, etc. (cf. [4,11,15,22,25,29,32,36]). More recently, efficient solutions were proposed for the plane [17] and for grids [16].

In many papers on gathering the agents are a priori assumed to have limited knowledge of the environment. Moreover, most papers supposed that an agent is not aware of the positions in the environment of other agents. In the deterministic settings, one of the central studied questions was feasibility of gathering or rendezvous, cf. [21,22,32], which most often led to some form of the symmetry breaking problem, see [29,31]. Surprisingly, when agents were equipped with GPS devices, knowledge of the agent's own position in the environment permitted executing very efficient rendezvous algorithms (see [14,15]).

Fault tolerance in mobile agent algorithms has also been extensively studied in the past, but the failures were more often related to the static elements of the environment (network nodes or links), cf. [26,30]. The faults of the mobile agents were studied for the problems of convergence [12], flocking [35], searching [19,20] or patrolling [18]. Faults or imperfections arriving to mobile agents performing gathering were investigated in [1,13,23,27,33]. Research in [13], [27] and [33] considered the gathering problem in the presence of inaccurate or faulty robot perception components. In [1] the initial positions of the collection is known to all robots, which operate in so called *look-compute-move* cycle. The feasibility of the problem, as a function of faulty robots, is investigated in [1] for crash and byzantine faults. In [23], the gathering problem is studied in an unknown graph environment and the feasibility question for byzantine faults in the strong and weak sense are investigated. The results of [23] depend on the knowledge of the upper bound on the size of the graph environment (or the absence of such knowledge).

In [9] and [10] the authors studied, similar to ours, the online rendezvous problem using GPS-equipped robots on a line, where some robots may turn out to be byzantine. However the robot movements along the line are much easier to analyze than the setting studied in the present paper. Indeed, in the case of a line, the robots move inside a corridor forcing robots to meet.

1.5 Notation

We will use $\mathcal{S}$ to refer to a general collection of any robots (reliable and/or byzantine) and use $\mathcal{N}$ ($\mathcal{F}$) to represent a set of reliable (byzantine) robots only. We will represent the cardinality of a set $\mathcal{S}$ as $|\mathcal{S}|$ and will

3

always use $n = |\mathcal{S}|$, and $f = |\mathcal{F}|$. We reserve the use of F for the upper bound on the number of byzantine robots in $\mathcal{S}$ (and, as such, it may be that $f \leq F$).

As we are dealing with robots in the plane we will use the term robot and point interchangeably. When it is required to refer to a particular robot / robots in a set we will use the capital letters A, B, and C. We use the capital letter D to refer to meeting points of robots.

We let the distance between any two points A and B be $|AB|$, and use $\overline{AB}$ to represent the directed line segment joining A and B. We will refer to the individual coordinates of a point using the subscripts x and y, e.g., $A = (A_x, A_y)$.

We define $\mathcal{MC}(\mathcal{S})$ as the minimum enclosing circle (MEC) of a set of points $\mathcal{S}$, and let $\text{Sup}[\mathcal{S}]$ be the supporting set of $\mathcal{MC}(\mathcal{S})$. It is a well known property that $2 \leq |\text{Sup}[\mathcal{S}]| \leq 3$ [8]. We further define the radius $\text{Radius}[\mathcal{S}]$ and $\text{Center}[\mathcal{S}]$ of $\mathcal{S}$ to be the radius and center of the MEC of $\mathcal{S}$ respectively.

Finally, we let $\mathcal{FVD}(\mathcal{S})$ represent the furthest-point Voronoi diagram (FVD) of the point set $\mathcal{S}$, and, for a point A in $\mathcal{S}$, we let $\mathcal{FVR}(A)$ be the cell / region in $\mathcal{FVD}(\mathcal{S})$ belonging to the point A. See [6] for a description of the properties of the FVD.

2 One byzantine robot

In this section we develop optimal algorithms for the case that there is only a single byzantine robot within the collection $\mathcal{S}$.

To do this we will need to consider subsets of $\mathcal{S}$ containing $n-1$ robots and we therefore introduce some convenient notation. We let $\mathcal{S}_i \subset \mathcal{S}$, $i \in [0, n-1]$ represent the n subsets of $n-1$ robots that can be formed from $\mathcal{S}$ and we define an ordering for the $\mathcal{S}_i$ in such a way that $\text{Radius}[\mathcal{S}_i] \leq \text{Radius}[\mathcal{S}_j] \ \forall \ j \geq i$. For the sake of brevity, we use $r_\mathcal{S} = \text{Radius}[\mathcal{S}]$ and $r_i = \text{Radius}[\mathcal{S}_i]$ for the remainder of the section.

We start with the following (trivial) lemma concerning the optimal meeting time of any set of robots in the plane,

Lemma 1. *The minimal time necessary to gather any set $\mathcal{S}$ of robots is $T_*(\mathcal{S}) = r_\mathcal{S}$.*

An immediate consequence of the above lemma is the following optimal algorithm for gathering a group of n reliable robots.

Algorithm 1 (Optimal $Gather(n, 0)$)

1: Set $D = \text{Center}[\mathcal{S}]$;
2: All robots in $\mathcal{S}$ move at full speed towards D ;
3: The algorithm terminates when the last robot in $\mathcal{S}$ reaches D ;

To get an idea of how different the problem is when we consider the presence of even a single byzantine robot, let us run the above algorithm on the two inputs depicted in Figure 2.
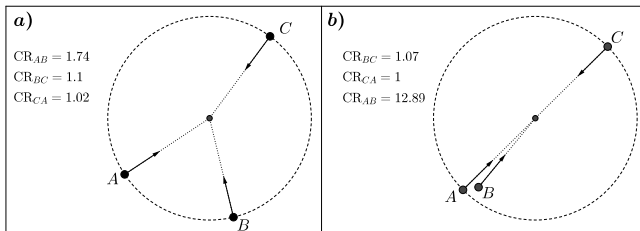

Fig. 2: Inputs for example analysis of competitive ratio. In both cases the robots A, B, and C move directly towards the center of the minimum enclosing circle of $\mathcal{S} = \{A, B, C\}$.
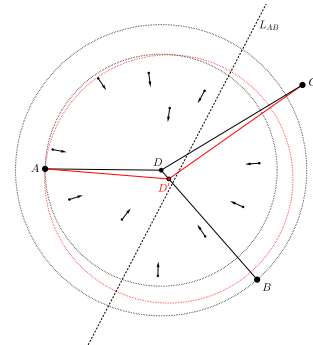

Fig. 3: Setup for the proof of Lemma 4.

For a given input $\mathcal{S} = \{A, B, C\}$ the adversary can choose at most one of the robots A, B, and C to be byzantine. We assume that they will do so in such a way as to maximize the competitive ratio of our algorithm. Which robot would they choose?

In the case **a)** the choice is not so obvious, and, indeed, the competitive ratios for all three possibilities are not very different. In the case **b)**, however, there *is* an obvious choice: the adversary would make C byzantine since the robots A and B were initially very close but travelled far before meeting.

This exercise, although simple, highlights an important observation – the "closest" robots should meet first. It turns out that, when $F = 1$, we can formalize this statement[6].

Lemma 2. *Consider an optimal algorithm $\overline{\mathcal{A}}$ solving the Gather(n, 1) problem for the input $\mathcal{S}$. Let $\mathcal{S}_i$ be the first group of $n - 1$ robots to meet. Then $\mathcal{S}_i = \mathcal{S}_0$, i.e. $\mathcal{S}_i$ is the group of $n - 1$ robots in $\mathcal{S}$ with the smallest enclosing circle.*

Proof. (Lemma 2) Assume we have an optimal algorithm $\overline{\mathcal{A}}$ such that $\mathcal{S}_i$ – the first group to gather using $\overline{\mathcal{A}}$ – is not the group with the smallest enclosing circle, i.e. $\mathcal{S}_i \neq \mathcal{S}$. In this case an adversary chooses $\mathcal{N} = \mathcal{S}_0$. Since the minimal time at which all robots can gather is $r_{\mathcal{S}}$, the competitive ratio of $\overline{\mathcal{A}}$ is $\geq r_{\mathcal{S}}/r_0$.

Now apply Algorithm 1 to solve this $Gather(n, 1)$ problem and observe that the competitive ratio of this algorithm is *equal* to $r_{\mathcal{S}}/r_0$. This implies that $\overline{\mathcal{A}}$ is, at best, as good as Algorithm 1. However, Algorithm 1 is not an optimal algorithm solving the $Gather(n, 1)$ problem. Thus, we must conclude that $\overline{\mathcal{A}}$ is not an optimal algorithm either – a contradiction. $\square$

So, we now know that we have to make the smallest group of $n - 1$ robots meet first. What choice does this leave the adversary? Well, naturally, they would choose the byzantine robots in such a way that the second-smallest group of $n - 1$ robots should have gathered. This observation leads us to the following:

Theorem 1. *The competitive ratio of any algorithm solving the Gather(n, 1) problem with input $\mathcal{S}$ is at least $r_{\mathcal{S}}/r_1$.*

Proof. (Theorem 1) Consider an algorithm $\mathcal{A}$ solving the $Gather(n, 1)$ problem with input $\mathcal{S}$. Let $\mathcal{S}_i$ be the first group of $n - 1$ robots to meet using $\mathcal{A}$, and let $\mathcal{S}_j = \mathcal{N}$ be the group of $n - 1$ reliable robots. Observe that an adversary can always choose to make $\mathcal{S}_i \neq \mathcal{S}_j$ such that, effectively, all n robots must meet before $\mathcal{A}$ terminates. Let the time at which this happens be T. Lemma 1 tells us that $r_{\mathcal{S}}$ is the minimum time necessary to gather all n robots and we thus have $T \geq r_{\mathcal{S}}$. The competitive ratio of $\mathcal{A}$ is therefore at least $\widehat{CR} \geq r_{\mathcal{S}}/r_j$, $j \neq i$. There are two cases to consider: $i = 0$ and the adversary chooses $j = 1$ such that $CR \geq \frac{r_{\mathcal{S}}}{r_1}$, or, $i \neq 0$ and the adversary chooses $j = 0$ such that $CR \geq \frac{r_{\mathcal{S}}}{r_0} \geq \frac{r_{\mathcal{S}}}{r_1}$. $\square$

At this point we can make a useful observation: an optimal gathering algorithm ends either at the moment the first group of robots meet or the moment all robots meet. Furthermore, at the moment of the first meeting, all robots are located at either one of only two positions. Thus, in an optimal algorithm, we must send these remaining two groups of robots directly towards each other. We can claim the following:

Lemma 3. *An optimal algorithm $\overline{\mathcal{A}}$ solving the Gather(n, 1) problem can be completely described by the single point D at which the first $n - 1$ robots gather.*

Corollary 1. *(Lemma 3) There is an optimal algorithm solving the Gather(n, 1) problem following the strategy given in Algorithm 2.*

Algorithm 2 (General $Gather(n, 1)$)

1: All n robots start moving at full speed towards some point D ;
2: **if** The first $n - 1$ robots to arrive at D are all reliable; **then**
3:      The algorithm terminates ;
4: **else**
5:      Let D' be the midpoint of D and the position of the single robot that has not yet arrived at D (at the time the first group of robots gather at D) ;
6:      All robots move at full speed towards D'. The algorithm terminates once they meet ;

---

[6] When $F > 1$ there are cases when this is not true.

Corollary 1 reduces the task of searching for an optimal algorithm to the conceptually simpler task of searching for some optimal meeting point D. The following lemma tells us how to find this point:

Lemma 4. *Consider an optimal algorithm $\overline{\mathcal{A}}$ solving the Gather(n, 1) problem for the input $\mathcal{S}$ parameterized by the point D. Let the group $\mathcal{S}_i$ represent the first group of $n-1$ robots to gather at the point D. Then the point D lies on the perpendicular bisector of the two robots in $\mathcal{S}_i$ furthest from D.*

Proof. (Lemma 4) Let A and B be the last two robots in $\mathcal{S}_i$ that reach D and let C be the single robot in $\mathcal{S}$ that is not contained in $\mathcal{S}_i$. We argue by contradiction and assume that the point D does not lie on the perpendicular bisector L_{AB} of A and B. Without loss of generality assume that A reaches D before B. This situation is depicted in Figure 3.

If we let the group of $n-1$ reliable robots be $\mathcal{S}_j = \mathcal{N}$ then the competitive ratio of $\overline{\mathcal{A}}$ is

$$\widehat{\mathrm{CR}}(\overline{\mathcal{A}}, \mathcal{S}) = \max_{S_j} \begin{cases} \frac{1}{r_i}|BD|, & \mathcal{S}_i = \mathcal{S}_j \\ \frac{1}{2r_j}(|BD| + |CD|), & \text{otherwise} \end{cases}.$$

Now consider the algorithm $\mathcal{A}'$ which replaces the point D in $\overline{\mathcal{A}}$ with the point D' on the segment $\overline{BD}$ located some small distance ϵ in the direction of L_{AB} such that A and B are still the last two robots to arrive at D' (and A arrives before B). The competitive ratio of this new algorithm is,

$$\widehat{\mathrm{CR}}(\mathcal{A}', \mathcal{S}) = \max_{S_j} \begin{cases} \frac{1}{r_i}|BD'|, & \mathcal{S}_i = \mathcal{S}_j \\ \frac{1}{2r_j}(|BD'| + |CD'|), & \text{otherwise} \end{cases}.$$

We claim that $\widehat{\mathrm{CR}}(\mathcal{A}', \mathcal{S}) < \widehat{\mathrm{CR}}(\overline{\mathcal{A}}, \mathcal{S})$. It is obvious that $|BD'| < |BD|$. Thus we need to show that $|BD'| + |CD'| < |BD| + |CD|$. Indeed, observe that $|BD| + |CD| = |BD'| + |DD'| + |CD| > |BD'| + |CD'|$ (triangle inequality). We can thus conclude that $\widehat{\mathrm{CR}}(\mathcal{A}', \mathcal{S}) < \widehat{\mathrm{CR}}(\overline{\mathcal{A}}, \mathcal{S})$ which is in contradiction to our assumption that $\overline{\mathcal{A}}$ is an optimal algorithm. □

As a last step we derive an expression for the competitive ratio of an optimal *Gather(n, 1)* algorithm.

Lemma 5. *An optimal algorithm following the strategy in Algorithm 2 solves the Gather(n, 1) problem for the input $\mathcal{S}$ with competitive ratio $\widehat{CR} = \max \left\{ \frac{|AD|}{r_0}, \frac{|AD|+|CD|}{2r_1} \right\}$ where A is one of the two points in $\mathcal{S}_0$ furthest from D and C is the point in $\mathcal{S}$ that is not in $\mathcal{S}_0$.*

Proof. (Lemma 5) Lemma 2 tells us that the first group of $n-1$ robots to gather is the group $\mathcal{S}_0$. Thus, if A is the point in $\mathcal{S}_0$ that is furthest from D, then, in the case that $\mathcal{S}_0 = \mathcal{N}$, the competitive ratio of the algorithm is $|AD|/r_0$.

If $\mathcal{S}_0 \neq \mathcal{N}$, then, if C is the single point in $\mathcal{S}$ that is not in $\mathcal{S}_0$, the algorithm terminates after time $(|AD| + |CD|)/2$. Thus, the overall competitive ratio of the algorithm is $\widehat{\mathrm{CR}}(D) = \max \left\{ \frac{|AD|}{r_0}, \frac{|AD|+|CD|}{2r_1} \right\}$ as required. □

We are now ready to present our main result:

Algorithm 3 (Optimal *Gather(n, 1)* point)

1: Set C as the single robot in $\mathcal{S}$ that is not in $\mathcal{S}_0$;
2: Determine the Furthest-point Voronoi diagram $\mathcal{FVD}(\mathcal{S}_0)$ of the point set $\mathcal{S}_0$;
3: Set $CR_{min} = \infty$, and $D_{min} = NULL$;
4: **for** each edge E in $\mathcal{FVD}[\mathcal{S}_0]$ **do**
5:     Set A and B as the two points such that the edge E separates $\mathcal{FVR}(A)$ and $\mathcal{FVR}(B)$;
6:     Determine the point D' on E that minimizes $\mathrm{CR}(D') = \max \left\{ \frac{|AD'|}{r_0}, \frac{|AD'|+|CD'|}{2r_1} \right\}$.
7:     **if** $\mathrm{CR}(D') < \mathrm{CR}_{min}$ **then**
8:         Set $D_{min} = D'$ and $\mathrm{CR}_{min} = \mathrm{CR}(D')$
    return D_{min};

Algorithm 4 (Optimal $Gather(n, 1)$)

1: The robots perform Algorithm 2 with the point D determined by Algorithm 3;

Theorem 2. *Algorithm 4 is an optimal algorithm solving the Gather(n, 1) problem with input $\mathcal{S}$. The complexity of the algorithm is $\mathcal{O}(n \log n)$.*

Proof. (Theorem 2) By Lemmas 2 and 4 we know that the optimal point D must lie on the perpendicular bisector of two points in $\mathcal{S}_0$ that are furthest from D. In Algorithm 3 we are choosing D to be on one of the edges of the FVD of $\mathcal{S}_0$. Thus, by construction, D does in fact lie on the perpendicular bisector of two points in $\mathcal{S}_0$ that are furthest from D.

The fact that the algorithm is optimal follows from Lemma 5 and the definition of an optimal algorithm given in Eq. (1).

The complexity of the algorithm is $\mathcal{O}(n \log n)$ since we need to determine the FVD of $\mathcal{S}_0$ which can be found optimally in $\mathcal{O}(n \log n)$ time [6]. Minimizing the quadratic equation given in Lemma 5 on a line segment can be done in constant time, and this needs to be done only once for each of the $\mathcal{O}(n)$ edges of the FVD. □

It does not seem likely that a closed form expression can be derived for the competitive ratio of Algorithm 4 for arbitrary n. However, in the boundary case that $n = 3$ and $F = 1$ this is possible. The complete solution of the $Gather(3, 1)$ is presented in the appendix and the results are reproduced below:

Theorem 3. *Algorithm 2 optimally solves the Gather(3, 1) problem with input $\triangle ABC$ of side lengths $a \leq b \leq c$ and respective angles $\alpha \geq \beta \geq \gamma$ if the point D is chosen such that $D_x = \frac{1}{2}[(B_x + C_x) + a \tan \phi(B_y - C_y)]$, $D_y = \frac{1}{2}[(B_y + C_y) + a \tan \phi(C_x - B_x)]$, and $\tan \phi = \tan \beta$ if $\tan \beta \leq \sin \gamma$, otherwise*

$$\tan \phi = \frac{2\sqrt{c^2 - (b-a)^2}}{\sqrt{(3b-a)^2 - c^2} + \sqrt{(b+a)^2 - c^2}}.$$

The competitive ratio of the algorithm equals c/b if $\tan \beta \leq \sin \gamma$, otherwise it is $1/\cos \phi$.

3 Bounded number of byzantine robots

We now consider instances of the $Gather(n, F)$ problem when the value of F is a small constant fraction of n. We give two algorithms corresponding to the cases that $F < \lceil \frac{n}{3} \rceil$, and $F < \lceil \frac{n}{2} \rceil$. In both cases we show that a small constant competitive ratio is attainable.

We start with the case that $F < \lceil \frac{n}{3} \rceil$. We have the following:

Theorem 4. *Consider the Gather(n, F) problem with input $\mathcal{S}$ and for any $F < \lceil \frac{n}{3} \rceil$. Then, there is a gathering algorithm solving this problem with competitive ratio at most 2. The complexity of the algorithm is $\mathcal{O}(n)$.*

Proof. (Theorem 4) We will make use of the centerpoint theorem (see [24] [Theorem 4.3]) which states that any finite set $\mathcal{S}$ of n points in $\mathbb{R}^d$ admits a point K (a centerpoint) such that any open half-space avoiding K contains at most $\lfloor \frac{dn}{d+1} \rfloor$ points of $\mathcal{S}$. In particular, for $d = 2$, this implies that we can always determine a K such that any line L through K partitions $\mathcal{S}$ into two sets each with at least $F < \lceil \frac{n}{3} \rceil$ robots. This result inspires the following algorithm,

---

Algorithm 5 (Move to centerpoint)

1: The robots compute a centerpoint K of the set $\mathcal{S}$ of robots;
2: All robots move directly towards K;
3: The algorithm terminates once the final reliable robot reaches K;

Consider the reliable robot A that is initially furthest away from the point K determined in Algorithm 5. Draw a line L through K perpendicular to the line segment $\overline{AK}$ (as done in Figure 4). Observe that, since K is a centerpoint, there are at least $\lceil \frac{n}{3} \rceil$ robots on either side of L. Furthermore, by assumption, F is strictly less than $\lceil \frac{n}{3} \rceil$ and we are thus guaranteed to have a reliable robot on either side of L. Consider any reliable robot B on the opposite side of L as A and note that the robot B is at least a distance $|AB| \geq |AK|$ away from the robot A. The competitive ratio of Algorithm 5 is therefore at most $\widehat{CR} \leq |AK|/(\frac{1}{2}|AB|) \leq 2$.

The complexity bound follows from the need to determine the centerpoint of the collection. The centerpoint of a set of n points can be determined in $\mathcal{O}(n)$ time using an algorithm by Jadhav [28].
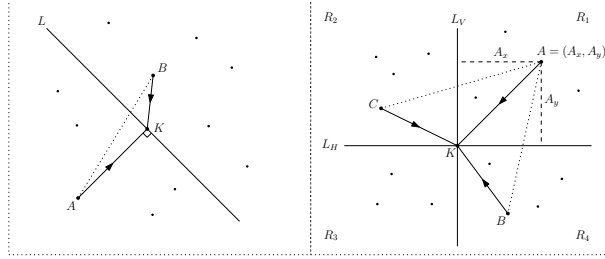

Fig. 4: Setup for the proofs of Theorem 4 (left) and Theorem 5 (right).

The centerpoint theorem applies generally to any d-dimensional space and we thus have the following corollary,

Corollary 2. *(Theorem 4) Consider the Gather(n, F) problem in $\mathbb{R}^d$ for for any $F < \lceil \frac{n}{d+1} \rceil$. Then, there exists a gathering algorithm with competitive ratio at most 2.*

Now consider the case that $F < \lceil \frac{n}{2} \rceil$. We claim the following:

Theorem 5. *Consider the Gather(n, F) problem with input $\mathcal{S}$ and for any $F < \lceil \frac{n}{2} \rceil$. Then, there is a gathering algorithm solving this problem with competitive ratio at most $2\sqrt{2}$. The complexity of the algorithm is $\mathcal{O}(n)$.*

Proof. (Theorem 5) The proof is based on the following algorithm,

---

Algorithm 6 (Move to intersection)

1: The robots compute a line L_H that partitions the robots into two disjoint sets each containing at least $\lceil \frac{n}{2} \rceil$ robots;
2: The robots compute a line L_V, perpendicular to L_H, that also partitions the robots into two disjoint sets each containing at least $\lceil \frac{n}{2} \rceil$ robots;
3: The robots move towards the point K that is the intersection of L_H and L_V.
4: The algorithm terminates once the final reliable robot reaches K;

First, we note that, in Algorithm 6, the existence of the lines L_H and L_V is ensured as a result of the ham-sandwich theorem (see [24] [Theorem 4.7]).

Now consider the four open regions R_1, R_2, R_3, and R_4 created by the intersection of L_H and L_V (as depicted in Figure 4). Note that, by assumption, we have $F < \lceil \frac{n}{2} \rceil$ and we are therefore guaranteed to have at least one reliable robot in each of the regions R_1 and R_3, or in each of the regions R_2 and R_4.

Consider the reliable robot A that is furthest from K and assume without loss of generality that A is located in the region R_1. If there is a reliable robot B in R_3 then we have $|AB| \geq |AK|$ which implies that $\widehat{CR} \leq |AK|/(\frac{1}{2}|AB|) \leq 2$. If there is not a reliable robot in R_3 then there must be reliable robots B and C in R_2 and R_4 respectively. Let $d = \max\{|AB|, |AC|\}$ and let us adopt a coordinate system such that $K = (0, 0)$ and $A = (A_x, A_y)$. Observe that $A_y \leq |AB| \leq d$ and $A_x \leq |AC| \leq d$. Thus, $|AK| = \sqrt{A_x^2 + A_y^2} \leq \sqrt{2}d$ and $\widehat{CR} \leq |AK|/(\frac{1}{2}d) \leq 2\sqrt{2}$.

The two lines L_H and L_V may be found in linear time by first choosing some line L' onto which we project the points in $\mathcal{S}$. We then set L_H as the line perpendicular to L' dividing the points on L' in half (i.e. we need to find the median, $\mathcal{O}(n)$ time [7]). To find L_V we repeat with L' replaced with L_H.

4 Arbitrary number of byzantine robots

In this section we consider algorithms that solve the $Gather(n, F)$ for any n and any F. We give two algorithms: the first, grid-rendezvous, is adapted from [14] and gives a constant competitive ratio independent of F. The second, shrinking-the-shortest-interval (SSI), gives a competitive ratio dependent on F.

4.1 Grid rendezvous

We start with the grid-rendezvous algorithm which is a direct application of Algorithm 3 in [14]. The algorithm was originally designed to solve the rendezvous problem of two robots unaware of the other's position (but sharing a common coordinate system).

The idea of the algorithm is to calculate a hierarchy of grids $\Pi = \{\pi_0, \pi_1, ...\}$ which partition the plane into non-overlapping cells. The robots then travel through a series of potential meeting points located at the centers of ever larger cells from successive grids in Π.

In detail, each π_i exactly partitions the plane into square cells of side length 2^i such that one of the cells in π_i, the central cell, has its center at the origin. In order for the partition to be exact each cell is defined to include its top and right edges, as well as its top-right vertex (in addition to its interior).

We can nearly apply Algorithm 3 as given in [14]. We only need to specify the finest grid division that will be used by the robots. Let d_ϵ be the size of this finest grid cell. We present (the slightly modified) Algorithm 3 from [14] below.

Algorithm 7 (Grid-rendezvous [14])

1: The robots choose a d_ϵ much smaller than the closest pair of robots in the set;
2: The robots compute the hierarchy of grids Π;
3: **repeat** for $i = 1, 2, 3...$ and for each robot in $\mathcal{S}$
4:     Set H equal to the cell of π_i containing your initial position p;
5:     Move to the center of H;
6:     Wait until $\sqrt{2} \cdot 2^{i-1}$ time has passed since the start of the current iteration;
7: **until** Gathering completed

The rendezvous time of the above algorithm is given by Corollary 9 in [14]. Using this time-bound we can state the following:

Theorem 6. *Consider the Gather(n, F) problem for the input $\mathcal{S}$. Assume that the robots A and B are the closest pair of robots in $\mathcal{S}$. Then the competitive ratio of Algorithm 7 is $\widehat{CR} \leq 2\sqrt{2}\left(16 + \frac{d_\epsilon}{|AB|}\right)$ where d_ϵ can be made as small as one chooses. The complexity of this algorithm is[7] $\mathcal{O}(n \log n)$.*

Proof. (Theorem 6) Choose any two robots I and J in $\mathcal{S}$ and assume that all distances are scaled such that $d_\epsilon = 1$. Then Corollary 9 from [14] tells us that the robots I and J gather in time $T \leq 16\sqrt{2} \cdot |IJ| + \sqrt{2}$.

Let $\mathcal{N}$ be the subset of reliable robots and consider the two robots A and B which are the most distant in $\mathcal{N}$. Then the minimal time necessary to gather the robots in $\mathcal{N}$ is at least $T_* \geq |AB|/2$. The competitive ratio of Algorithm 7 is therefore $\widehat{CR} = \frac{T}{T_*} \leq \frac{16\sqrt{2}|AB|+\sqrt{2}}{\frac{1}{2}|AB|} \leq 2\sqrt{2}\left(16 + \frac{1}{|AB|}\right)$. In the worst case, $|\mathcal{N}| = 2$ and the robots A and B were the closest pair in $\mathcal{S}$.

The complexity of the algorithm is $\mathcal{O}(n \log n)$ as one needs to find the closest pair of points in $\mathcal{S}$ in order to determine a satisfactory grid division d_ϵ. The closest pair in a set of n points can be found optimally in $\mathcal{O}(n \log n)$ time [5]. □

---
[7] The complexity of the algorithm is entirely due to the determination of d_ϵ.

4.2 Shrink-shortest-interval

Consider the following algorithm, generalized from Algorithm 3 in [10]:

Algorithm 8 (Shrink-shortest-interval)

1: **repeat**
2:     Determine the two closest robots A and B in $\mathcal{S}$ that are not at the same position;
3:     Set D as the midpoint of A and B;
4:     Set $d = |AB|/2$.
5:     All robots move a distance d towards D;
6: **until** All robots in $\mathcal{N}$ gather.

We claim the following:

Theorem 7. *Algorithm 8 solves the Gather(n, F) problem for the input $\mathcal{S}$ with competitive ratio at most $F + 2$. The complexity of the algorithm is $\mathcal{O}(n^2 \log n)$.*

To prove this we will need the following lemma:

Lemma 6. *Consider any point D and set of points $\mathcal{S}$ such that $A \in \mathcal{S}$ is the closest point to D, and $C \in Sup[\mathcal{S}]$ is the furthest point from D. Let $\mathcal{S}'$ be the positions of the points in $\mathcal{S}$ after moving them a distance $d \leq |AD|$ towards the point D. Then,*

$$Radius[\mathcal{S}'] \leq \begin{cases} Radius[\mathcal{S}] - d/2, & D \in \mathcal{MC}(\mathcal{S}) \\ Radius[\mathcal{S}], & otherwise \end{cases}.$$

Proof. (Lemma 6) Let K and $r_\mathcal{S}$ be the center and radius of $\mathcal{MC}(\mathcal{S})$ respectively and adopt a coordinate system which places K at the origin and oriented such that the line segment $\overline{KD}$ is along the positive x-axis. Then $D = (D_x, 0)$ and, for any point C on the border of $\mathcal{MC}(\mathcal{S})$ we have $C = r_\mathcal{S}(\cos\theta, \sin\theta)$ where θ is the angle between $\overline{KC}$ and the x-axis.

Define the point C' as the point obtained by moving the point C a distance $d \leq |CD|$ towards D. We can write $C'_x = r_\mathcal{S}\cos\theta + \frac{d}{|CD|}(D_x - r_\mathcal{S}\cos\theta) = r_\mathcal{S}\cos\theta\left(1 - \frac{d}{|CD|}\right) + \frac{d \cdot D_x}{|CD|}$, and, $C'_y = r_\mathcal{S}\sin\theta - \frac{d}{|CD|} \cdot r_\mathcal{S}\sin\theta = r_\mathcal{S}\sin\theta\left(1 - \frac{d}{|CD|}\right)$. Observe that the equations for C'_x, and C'_y describe a parametric curve that is completely contained within a circle of radius $r_\mathcal{S}(1 - d/|CD|) \leq r_\mathcal{S}$. Thus, for any point D, we can conclude that $Radius[\mathcal{S}'] \leq r_\mathcal{S}(1 - d/|CD|) \leq r_\mathcal{S}$.

Now consider the case that D is inside $\mathcal{MC}(\mathcal{S})$. In this case $|CD| \leq 2r_\mathcal{S}$ such that the curve defined by C'_x, and C'_y is completely contained within a circle of radius $r_\mathcal{S}(1 - d/|CD|) \leq (r_\mathcal{S} - d/2)$. □

Proof. (Theorem 7) Consider the $Gather(n, F)$ problem for the input $\mathcal{S}$, let $\mathcal{N}$ be the subset of $\mathcal{S}$ that contains only reliable robots, and let f be the (actual) number of byzantine robots in $\mathcal{S}$.

Let $\mathcal{S}^{(i)}$ and $\mathcal{N}^{(i)}$ represent the unique positions of the robots in $\mathcal{S}$ and $\mathcal{N}$ after the i^{th} iteration of the algorithm, and let $r_i = Radius[\mathcal{N}^{(i)}]$. We also let D_i be the midpoint and d_i be half the distance between the closest pair of points in $\mathcal{S}^{(i)}$. Finally, set $C_i \in Sup[\mathcal{N}^{(i)}]$ be the furthest point from D_i.

Now, if in the i^{th} iteration the midpoint D_i lies within $\mathcal{MC}(\mathcal{N}^{(i)})$ then by Lemma 6 we have $r_{i+1} \leq r_i - d_i/2$. If we assume that their are m iterations of this kind then the time needed to complete these iterations is at most $T_m \leq \sum_{i=0}^{m} d_i \leq 2\sum_{i=0}^{m}(r_i - r_{i+1})$. However, observe that $\sum_{i=0}^{m} r_i = r_0 + \sum_{i=1}^{m} r_i = r_0 + \sum_{i=0}^{m-1} r_{i+1}$ such that $T_m \leq 2r_0 - r_{m+1} \leq 2r_0$.

If D_i does not lie within $\mathcal{MC}(\mathcal{N}^{(i)})$, then we can only say that $r_{i+1} \leq r_i(1 - d_i/|C_iD_i|) \leq r_i$. However, observe that Algorithm 8 always gathers the two closest robots in $\mathcal{S}^{(i)}$ and we know that there is at least one pair of robots in $\mathcal{N}^{(i)}$ with separation no greater than $2r_i$. This tells us that $d_i \leq r_i$. Furthermore, since all reliable robots are, by definition, within $\mathcal{MC}(\mathcal{N})$, it is impossible for D_i to simultaneously be: a) the midpoint of two reliable robots, and, b) lie outside of $\mathcal{MC}(\mathcal{N})$. This implies that this type of iteration can

occur at most f times (as it reduces the number of byzantine robots by one each time it occurs). Thus, the time needed to complete these iterations is at most $T_f = f \cdot r_0$.

Combining T_m and T_f gives us a bound on the total time necessary to complete the algorithm. We get $T \le T_m + T_f = fr_0 + 2r_0 = (f+2)r_0$. The bound on the competitive ratio follows from the fact that $f \le F$, and $r_0 = \text{Radius}[\mathcal{N}]$ is the minimal time necessary to gather the robots in $\mathcal{N}$.

The complexity bound follows from the fact that we need to determine the closest pair of points $\mathcal{O}(n)$ times. $\qquad\square$

In the case that we have no knowledge of the number of byzantine robots in our collection (i.e. $F = n-2$) the algorithm has a worst-case bound on the competitive ratio of n. This reflects the fact that an adversary, if allowed, would always choose $f = F$ robots in $\mathcal{S}$ to be byzantine. It is worth noting, however, that it was not necessary to know F in the proof of Theorem 7 and thus the algorithm has a competitive ratio that is bounded by the actual number of byzantine robots in $\mathcal{S}$. That is, for a particular instance $\mathcal{N} \subseteq \mathcal{S}$ such that $f = |\mathcal{S}| - |\mathcal{N}|$ we have $\text{CR}(\mathcal{N}) \le f + 2 \le F + 2$.

5 Conclusion

In this paper we analyzed the gathering problem for $n > 2$ robots in the plane at most F of which, $F \le n-2$, are byzantine. The robots were equipped with GPS and they could communicate their positions to a central authority. Several algorithms were designed with competitive ratio depending on the number of byzantine robots and the knowledge available to the robots.

In addition to improving the competitive ratio and/or complexity of our algorithms, several interesting open problems remain. In particular, one could consider models that allow the robots to communicate/exchange their positions at any time during the gathering process. Additionally, it would be interesting to consider robot gathering (in the presence of byzantine robots) under local (limited) communication range.

References

1. N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM Journal on Computing*, 36(1):56–82, 2006.
2. S. Alpern. The rendezvous search problem. *SIAM Journal on Control and Optimization*, 33(3):673–683, 1995.
3. S. Alpern. Rendezvous search: A personal perspective. *Operations Research*, 50(5):772–795, 2002.
4. S. Alpern and S. Gal. *The theory of search games and rendezvous*, volume 55. Springer, 2003.
5. Jon Louis Bentley. Multidimensional divide-and-conquer. *Commun. ACM*, 23(4):214–229, April 1980.
6. Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, 3rd edition, 2008.
7. Manuel Blum, Robert W. Floyd, Vaughan Pratt, Ronald L. Rivest, and Robert E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448 – 461, 1973.
8. George Chrystal. On the problem to construct the minimum circle enclosing n given points in the plane. *Proceedings of the Edinburgh Mathematical Society*, 3:30–33, 1885.
9. H. Chuangpishit, J. Czyzowicz, R. Killick, E. Kranakis, and D. Krizanc. Optimal rendezvous on a line by location-aware robots in the presence of spies. (submitted), 2017.
10. H. Chuangpishit, J. Czyzowicz, E. Kranakis, and D. Krizanc. Rendezvous on a line of faulty, location-aware robots. In *Proceedings 13th International Symposium on Algorithms and Experiments for Wireless Networks*, Vienna, Austria, 2017. Springer, LNCS.
11. M. Cieliebak, P. Flocchini, P. Prencipe, and N. Santoro. Solving the robots gathering problem. In *Automata, Languages and Programming, 30th International Colloquium, ICALP 2003. Proceedings*, pages 1181–1196, Eindhoven, The Netherlands, June 30 - July 4, 2003. Springer, LNCS.
12. R. Cohen and D. Peleg. Convergence properties of the gravitational algorithm in asynchronous robot systems. *SIAM J. Comput.*, 34(6):1516–1528, 2005.
13. R. Cohen and D. Peleg. Convergence of autonomous mobile robots with inaccurate sensors and movements. *SIAM Journal on Computing*, 38(1):276–302, 2008.

14. A. Collins, J. Czyzowicz, L. Gasieniec, A. Kosowski, and R. Martin. Synchronous rendezvous for location-aware agents. In *International Symposium on Distributed Computing*, pages 447–459, Rome, Italy, September 20-22, 2011. Springer, LNCS.

15. A. Collins, J. Czyzowicz, L. Gasieniec, and A. Labourel. Tell me where I am so I can meet you sooner. In *International Colloquium on Automata, Languages, and Programming*, pages 502–514, Bordeaux, France, July 6-10, 2010. Springer, LNCS.

16. A. Cord-Landwehr, M. Fischer, D. Jung, and F. Meyer auf der Heide. Asymptotically optimal gathering on a grid. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2016*, pages 301–312, Asilomar State Beach/Pacific Grove, CA, USA,, July 11-13, 2016. ACM.

17. P. Courtieu, L. Rieg, S. Tixeuil, and X. Urbain. Certified universal gathering in R^2 for oblivious mobile robots. In *Distributed Computing - 30th International Symposium, DISC 2016, Proceedings*, pages 187–200, Paris, France, September 27-29, 2016. Springer, LNCS.

18. J. Czyzowicz, L. Gasieniec, A. Kosowski, E. Kranakis, D. Krizanc, and N. Taleb. When patrolmen become corrupted: Monitoring a graph using faulty mobile robots. In *Algorithms and Computation - 26th International Symposium, ISAAC 2015, Proceedings*, pages 343–354, Nagoya, Japan, December 9-11, 2015. Springer, LNCS.

19. J. Czyzowicz, K. Georgiou, E. Kranakis, D. Krizanc, L. Narayanan, J. Opatrny, and S. M. Shende. Search on a line by byzantine robots. In *27th International Symposium on Algorithms and Computation, ISAAC 2016*, pages 27:1–27:12, Sydney, Australia, December 12-14, 2016. Springer, LNCS.

20. J. Czyzowicz, E. Kranakis, D. Krizanc, L. Narayanan, and J. Opatrny. Search on a line with faulty robots. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016*, pages 405–414, Chicago, IL, USA, July 25-28, 2016. ACM.

21. S. Das, P. Flocchini, N. Santoro, and M. Yamashita. On the computational power of oblivious robots: forming a series of geometric patterns. In *Proceedings of the 29th PODC*, pages 267–276, Zurich, Switzerland, July 25-28, 2010. ACM.

22. G. De Marco, L. Gargano, E. Kranakis, D. Krizanc, A. Pelc, and U. Vaccaro. Asynchronous deterministic rendezvous in graphs. *Theoretical Computer Science*, 355(3):315–326, 2006.

23. Y. Dieudonné, A. Pelc, and D. Peleg. Gathering despite mischief. *ACM Transactions on Algorithms (TALG)*, 11(1):1, 2014.

24. H. Edelsbrunner. *Algorithms in combinatorial geometry*, volume 10. Springer Science & Business Media, 2012.

25. P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer. Gathering of asynchronous oblivious robots with limited visibility. In *STACS 2001, 18th Annual Symposium on Theoretical Aspects of Computer Science, Proceedings*, pages 247–258, Dresden, Germany, February 15-17, 2001. Springer, LNCS.

26. J. Hromkovič, R. Klasing, B. Monien, and R. Peine. Dissemination of information in interconnection networks (broadcasting & gossiping). In *Combinatorial network theory*, pages 125–212. Springer, 1996.

27. T. Izumi, S. Souissi, Y. Katayama, N. Inuzuka, X. Défago, K. Wada, and M. Yamashita. The gathering problem for two oblivious robots with unreliable compasses. *SIAM Journal on Computing*, 41(1):26–46, 2012.

28. S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. *Discrete & Computational Geometry*, 12(3):291–312, Sep 1994.

29. E. Kranakis, D. Krizanc, and S. Rajsbaum. Mobile agent rendezvous: A survey. In *Structural Information and Communication Complexity, 13th International Colloquium, SIROCCO 2006, Proceedings*, pages 1–9, Chester, UK,, July 2-5, 2006. Springer, LNCS.

30. N. A. Lynch. *Distributed algorithms*. Morgan Kaufmann, 1996.

31. A. Pelc. Deterministic rendezvous in networks: Survey of models and results. In *Proceedings of 29th DISC*, pages 1–15, Rome, Italy, September, 20-22, 2011. Springer, LNCS.

32. G. Prencipe. Impossibility of gathering by a set of autonomous mobile robots. *Theor. Comput. Sci.*, 384(2-3):222–231, 2007.

33. S. Souissi, X. Défago, and M. Yamashita. Gathering asynchronous mobile robots with inaccurate compasses. In *Principles of Distributed Systems (OPODIS)*, pages 333–349, Bordeaux, France, 12-15 December, 2006. Springer, LNCS.

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

35. Y. Yang, S. Souissi, X. Défago, and M. Takizawa. Fault-tolerant flocking for a group of autonomous mobile robots. *Journal of Systems and Software*, 84(1):29–36, 2011.

36. X. Yu and M. Yung. Agent rendezvous: A dynamic symmetry-breaking problem. In *ICALP*, pages 610–621, Kyoto, Japan, 6-10 July, 1996. Springer, LNCS.

A Gathering three robots

Observe that we can describe an instance of the $Gather(3, 1)$ problem by the triangle $\triangle ABC$ whose vertices specify the initial positions of the three robots for which gathering should occur. Thus, throughout this section, we let $a = |BC|$, $b = |AC|$, and $c = |AB|$ be the side lengths of $\triangle ABC$, and set the angles $\beta = \angle ABC$, and $\gamma = \angle BCA$. Without loss of generality we assume that $a \leq b \leq c$.

With this description of the problem Lemmas 2, 4, 5, and Theorem 1 take the following simple forms:

Corollary 3. *(Lemma 2 and Lemma 4) An optimal algorithm solving the Gather(3, 1) problem with input $\triangle ABC$ has the robots B and C meet first at some point on their perpendicular bisector.*

Corollary 4. *(Theorem 1 and Lemma 5) An optimal algorithm solving the Gather(3, 1) problem with input $\triangle ABC$ has competitive ratio*

$$\widehat{CR} = \max \left\{ \frac{2|BD|}{a}, \ \frac{|AD| + |BD|}{b} \right\} \tag{2}$$

and this is at least c/b.

At this point we could simply apply Algorithm 4 to determine the optimal point for this problem, however, in order to derive a closed form expression for the point D we take a different approach. We claim the following:

Theorem 3. *Algorithm 2 optimally solves the Gather(3, 1) problem with input $\triangle ABC$ of side lengths $a \leq b \leq c$ and respective angles $\alpha \geq \beta \geq \gamma$ if the point D is chosen such that $D_x = \frac{1}{2}[(B_x + C_x) + a \tan \phi (B_y - C_y)]$, $D_y = \frac{1}{2}[(B_y + C_y) + a \tan \phi (C_x - B_x)]$, and $\tan \phi = \tan \beta$ if $\tan \beta \leq \sin \gamma$, otherwise*

$$\tan \phi = \frac{2\sqrt{c^2 - (b-a)^2}}{\sqrt{(3b-a)^2 - c^2} + \sqrt{(b+a)^2 - c^2}}.$$

The competitive ratio of the algorithm equals c/b if $\tan \beta \leq \sin \gamma$, otherwise it is $1/\cos \phi$.

Proof. (Theorem 3) First consider the case that $\tan \beta \leq \sin \gamma$ (see Figure 5, top). In this case D lies on the edge $\overline{AB}$ and we have $2|BD|/a = 1/\cos \beta$ and $|BD| + |AD| = |AB| = c$. Thus, by Eq. (2), we have, $\widehat{CR} = \max \left\{ \frac{1}{\cos \beta}, \ \frac{c}{b} \right\}$. We claim that the condition $\tan \beta \leq \sin \gamma$ implies that $1/\cos \beta \leq c/b$. Indeed, observe that $\tan \beta \leq \sin \gamma$ which we can rewrite as $\frac{1}{\cos \beta} \leq \frac{\sin \gamma}{\sin \beta} = \frac{c}{b}$ (where we have invoked the sine law in the last step).
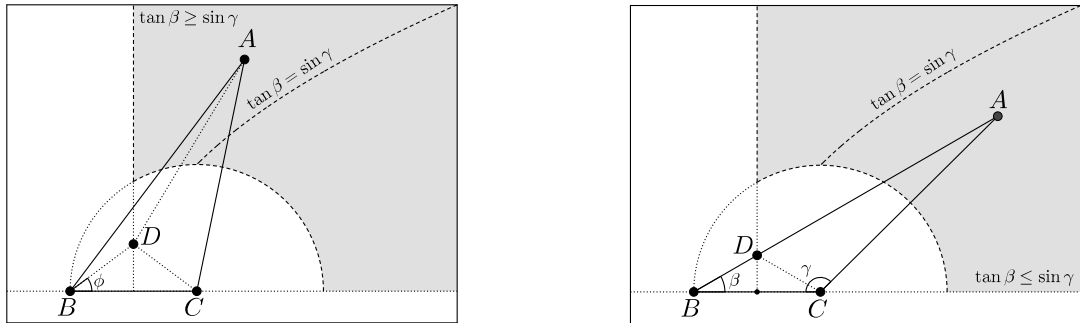

Fig. 5: Setup for the proof of Theorem 3. The shaded gray region indicates those positions of A such that $a \leq b \leq c$. Right: $\tan \beta \leq \sin \gamma$. Left: $\tan \beta \geq \sin \gamma$

Now consider the case that $\tan\beta \le \sin\gamma$ (see Figure 5, bottom) and define $\mathrm{CR}_A = \frac{2|BD|}{a}$, and $\mathrm{CR}_B = \frac{|BD|+|AD|}{b}$ such that $\widehat{\mathrm{CR}} = \max\{\mathrm{CR}_A,\ \mathrm{CR}_B\}$. We note that, over the interval $[0,\beta]$, CR_B is a monotone continuous decreasing function of ϕ and CR_A is a monotone continuous increasing function of ϕ. Furthermore, $CR_A \ge CR_B$ when $\phi = \beta$ and $CR_A = 1 \le CR_B$ when $\phi = 0$. We can thus conclude that the optimal competitive ratio is given when $CR_A = CR_B$, and the optimal point D is determined from the following:

$$\frac{2|BD|}{a} = \frac{|BD|+|AD|}{b}. \tag{3}$$

Now, let us choose a coordinate system such that the midpoint of B and C is at the origin, and the positive y-axis lies along the perpendicular bisector of B and C. In this coordinate system $B = (-a/2, 0)$, $C = (a/2, 0)$, $A = (A_x, A_y)$, and $D = (0, d)$ where we have introduced the parameter d. Thus, we can write $|BD| = \sqrt{\frac{1}{4}a^2 + d^2}$, and $|AD| = \sqrt{\left(A_x - \frac{a}{2}\right)^2 + A_y^2}$. Plugging these into Eq. (3) gives us (after some manipulation) the following quadratic equation $4b(b-a)d^2 + 2a\Delta d - \frac{a^2}{2}[c^2 - (b-a)^2] = 0$ where Δ is the area of $\triangle ABC$. Solving for d we get (after some manipulation),

$$d = \frac{a[c^2 - (b-a)^2]}{\sqrt{4\Delta^2 + 2b(b-a)(c^2 - (b-a)^2)} + 2\Delta}.$$

Finally, by applying Heron's formula for the area of a triangle, and noting that $d = \frac{a}{2}\tan\phi$, we get our final result:

$$\tan\phi = \frac{2\sqrt{c^2 - (b-a)^2}}{\sqrt{(3b-a)^2 - c^2} + \sqrt{(b+a)^2 - c^2}}.$$

□

It is interesting to note that in some cases ($\tan\beta \le \sin\gamma$) it was possible to achieve the lower-bound given in Theorem 1. This turns out to be true for the general case as well, and, in fact, it is possible to specify under which conditions this occurs:

Lemma 7. *Consider an input $\mathcal{S}$ such that $|Sup[\mathcal{S}]| = 2$. Let L be the line segment defined by the two points in $Sup[\mathcal{S}]$, and let $\mathcal{P}$ be the convex region defined by the intersection of all circles with centers given by the points in $\mathcal{S}_0$ and radii given by $r_0 \cdot r_\mathcal{S}/r_1$. Then, if L intersects with $\mathcal{P}$, the competitive ratio of Algorithm 4 is $\widehat{CR} = r_\mathcal{S}/r_1$.*

Proof. (Lemma 7) Define the convex region $\mathcal{P}$ as above and observe that $\mathcal{P} \ne \emptyset$ since, at minimum, it contains the center of $\mathcal{MC}(\mathcal{S}_0)$.

Now, we let A in $\mathcal{S}_0$ and $C \in \mathcal{S}_1$ be the two points that are also in $Sup[\mathcal{S}]$. Observe that for any point D on the segment $L = \overline{AC}$ we have $|AD| + |CD| = |AC| = 2r_\mathcal{S}$. Thus, if A is (one of) the furthest point(s) from D, the algorithm terminates in at most $r_\mathcal{S}$ time.

Now observe that for any point $D \in \mathcal{P}$ the distance between D and any point in $\mathcal{S}_0$ is at most q. In particular, for the point A, we have $|AD| \le q$.

Thus, if it happens that: a) $\mathcal{P}$ and L intersect, and, b) we can choose D in $\mathcal{P} \cap L$ in such a way that A is (one of) the furthest point(s) from D, then, the competitive ratio of Algorithm 2 is

$$\widehat{\mathrm{CR}} = \max\left\{\frac{|AD|}{r_0},\ \frac{|AD|+|CD|}{2r_1}\right\} \le \max\left\{\frac{q}{r_0},\ \frac{2r_\mathcal{S}}{2r_1}\right\} \le \max\left\{\frac{r_\mathcal{S}}{r_1},\ \frac{r_\mathcal{S}}{r_1}\right\} = \frac{r_\mathcal{S}}{r_1}.$$

Now, observe that there must be a point in $\mathcal{P} \cap L$ such that A is furthest from it. One such point, for example, is the point of intersection of the segment L and the circle of radius q centered on A. In order to get the optimal point we simply need to choose the D that minimizes the quadratic given in Lemma 5 on the portion of L contained in the region of $\mathcal{FVD}(\mathcal{S}_0)$ belonging to the point A. □