

Experimental Analysis of the Reynolds Flocking Model *

J.D. Eversham[†], V.F. Ruiz[‡]

Cybernetics, School of Systems Engineering
University of Reading, Reading, UK

Received 2011/11/09

Accepted 2012/01/17

Abstract

The classic Reynolds flocking model is formally analysed, with results presented and discussed. Flocking behaviour was investigated through the development of two measurements of flocking, flock area and polarisation, with a view to applying the findings to robotic applications. Experiments varying the flocking simulation parameters individually and simultaneously provide new insight into the control of flock behaviour.

Keywords

flocking · aggregation · traffic · control · swarm robotics · behaviour

1. Introduction

Flocking is the coherent and collective emergent motion of a number of birds, similar to motions observed in other animals, and has been investigated since the early 1970s [1]. Birds have been observed to apparently behave autonomously, moving together with high co-ordination in cohesive groups, which may contain a few individual birds to hundreds. The flocking models introduced by Reynolds [2] provided a method of simulating the appearance of bird flocking on a computer, through the use of simple rules. Such modelling has been used in films as *The Lion King* and *Batman Returns* to animate wildebeest and bats respectively. However, there have been limited other practical applications derived from flocking research. The main practical application resulting from research into flocking has been the optimisation technique Particle Swarm Optimisation [3].

Road traffic is another emergent phenomena, occurring when multiple vehicles travel in the same direction, impeding commuters and freight distribution on a daily basis. The cost of increasing road infrastructure is expensive, both from a financial consideration, but also through introducing additional delays to road users. Further, the Braess paradox shows that additional road infrastructure does not necessarily produce an improvement in traffic conditions [4].

An approach to decreasing traffic issues is to automate the act of driving, through the introduction of Intelligent Transport Systems. By transferring vehicle control from humans to machines, advanced road utilisation would be possible: for example, automated vehicles should be able to travel with smaller intra-vehicle spacing than humans can consistently and safely achieve. Through such developments, the density

of road traffic will be increased, hence allowing greater flow through road infrastructure.

Recent developments in robotics have increased the focus on autonomous robotic agents: for example, the autonomous transport vehicles that have participated in DARPA's Grand and Urban Challenges [5, 6]. Following from the development of autonomous individual agents, there is increased interest in using multiple, coordinated autonomous agents in a controlled and productive method. For example, multiple robots have been used to partially automate warehouse management [7]. Robot "swarms" are an area of active research, for use in domains including search-and-rescue, such as in the treacherous conditions following an earthquake [8] or gas leak [9–11], or for mining tasks [12]. These areas of investigation have been enabled by continued developments in computer speeds and decreasing costs of components, making research into autonomous robotics more accessible.

Consistent through these application areas is the use of multiple robots working together to achieve a common goal. It is considered that the application of flocking rules to the aforementioned application areas will support sub-goals including locomotion and area search. As such, it is further considered imperative to fully understand the mechanics of flocking models, so that they could be efficiently implemented in applications using multiple robotic agents.

The focus of this paper is hence to develop the required detailed understanding of how the Reynolds flocking model operates, including how the model is affected by modifications to the parameters involved, as this analysis appears to be lacking from the literature. New insight into how the flocking parameters impact the flock's overall behaviour could then be used in applications. For example, modifying the group's motion pattern in search-and-rescue or reconnaissance type tasks.

The rest of the paper is organised as follows. Section 2 presents the rules used in flocking models. In Section 3, the experiments performed in simulation are presented, and Section 4 gives and discusses the results obtained. Additional considerations ahead of applying the findings to application domains are presented in Section 5. The conclusions from the simulations are given in Section 6.

*This paper has been partially published in the IEEE 9th International Conference on Cybernetic Intelligent Systems, 2010, Reading, UK

[†]E-mail: j.d.eversham@reading.ac.uk

[‡]E-mail: v.f.ruiz@reading.ac.uk

2. Flocking Rules

This section outlines the flocking rules used in the Reynolds model, with a brief description of how they are implemented. In what follows, let b represent the individual bird under consideration and f_i another bird i in the flock f . Bird b has coordinates (b_x, b_y) and bird i has coordinates (f_{ix}, f_{iy}) .

The following equations will be represented in reference to the x axis: changing the index to y will give the respective result for the other direction.

The rules rely on the calculation of distance. We define the distance between the bird b and a bird i in the flock f , for birds $i = 0 \dots n$, by:

$$d_{i,x} = b_x - f_{ix} \quad (1)$$

Each of the three flocking rules is performed for all birds in the flock, in the order presented, to calculate each bird's new velocity, v_{t+1} . Once the three rules have been applied, the positions of all birds in the flock are updated. The rules are applied and the birds' positions updated in each time step of the simulation.

A number of the flocking rules make use of scaling parameters, which are detailed in Tables 1 and 2. These scaling parameters are used to adjust the effect each flocking rule has on the simulation.

2.1. Move Closer

For all birds in the flock, other than the individual bird under consideration, b :

Calculate the average distance, $|\bar{a}|$, between bird b and each other bird i in the flock, f :

$$|\bar{a}_x| = \sum_{i=0}^n \frac{(b_x - f_{ix})}{n}, \quad (b \neq f_i) \quad (2)$$

Divide the average distance by the Move Closer scaling factor, s_c , giving the scaled distance, s_d :

$$s_d = \frac{|\bar{a}_x|}{s_c} \quad (3)$$

Subtract the scaled distance from bird b 's velocity at time t , giving b 's new velocity at time $t + 1$:

$$b_{vx,t+1} = b_{vx,t} - s_d \quad (4)$$

2.2. Move With

Calculate the average flock velocity:

$$\bar{v}_x = \sum_{i=0}^n \frac{v_{ix}}{n} \quad (5)$$

Divide the average velocity by the Move With scaling factor, s_w , to produce a scaled average velocity, $s_{\bar{v}_x}$:

$$s_{\bar{v}_x} = \frac{\bar{v}_x}{s_w} \quad (6)$$

Add the scaled average velocity to bird b 's initial velocity, giving bird b 's final velocity:

$$b_{vx,t+1} = b_{vx,t} + s_{\bar{v}_x} \quad (7)$$

2.3. Move Away

If the distance between two birds given by 1 is less than a minimum distance, d_{min} , the birds are evaluated as being too close, and need to move apart:

$$b_x - f_{ix} < d_{min} \quad (8)$$

The set of birds that are shown to be too close are referred to with the j index in the following equations, to demarcate them from birds i in the flock, where birds $f_j \in f_i$.

Calculate the axis difference root, ad_x :

$$ad_{jx} = \sqrt{d_{min}^2 - d_{jx}^2} \quad (9)$$

Sum the resulting axis difference roots, then divide by the Move Away scaling factor s_a , giving the scaled summed axis difference root, S_x :

$$S_x = \frac{\sum_{i=0}^n ad_{jx}}{s_a} \quad (10)$$

Subtract S_x from the bird under consideration b 's velocity:

$$b_{vx,t+1} = b_{vx,t} - S_x \quad (11)$$

2.4. Update the Bird Position

Calculate the magnitude of bird b 's velocity, $b_{|v|}$:

$$b_{|v|} = \sqrt{b_{vx}^2 + b_{vy}^2} \quad (12)$$

If the magnitude of bird b 's velocity, $b_{|v|}$, exceeds a simulation defined maximum velocity, v_{max} , then the velocity to be applied will need to be reduced. This is to help ensure that birds cannot move too far in one time step.

If the magnitude of bird b 's velocity, $b_{|v|}$, does not exceed v_{max} , then 13 and 14 do not apply, and 15 is immediately applied.

Divide the velocity magnitude, $|v|$, by the defined maximum velocity v_{max} , giving the velocity scaling factor s_v :

$$s_v = \frac{|v|}{v_{max}} \quad (13)$$

Divide bird b 's velocity by the velocity scaling factor s_v , giving a scaled velocity, s_{vx} :

$$s_{vx} = \frac{b_{vx}}{s_v} \quad (14)$$

Update bird b 's position by adding either the scaled velocity s_{vx} or calculated velocity b_{vx} , depending on the value of $b_{|v|}$:

$$\begin{aligned} b_{x,t+1} &= b_{x,t} + s_{vx} \} \text{ if } b_{|v|} > v_{max} \\ b_{x,t+1} &= b_{x,t} + b_{vx} \} \text{ if } b_{|v|} \leq v_{max} \end{aligned} \quad (15)$$

3. Experimentation

3.1. Experimental setup

The parameters considered for experimentation include the scaling factors affecting the move closer, move with and move away flocking rules; the number of birds in the simulation; the maximum velocity that a bird can move in one time step; and the minimum distance that should separate two birds.

A two dimensional flocking simulation based on the original model by Reynolds [2] was developed. Values of experimental settings were selected through linearly changing the base values in the model.

Experiments were designed to understand how flocking parameters affect the flock size and polarisation (refer to the Metrics section for commentary on these measurements). These two measurements were selected as it is considered they are relevant to applications involving robotic agents deployed in two-dimensional settings. For example, in navigating and searching an area for a given target, it may be required to alter the group's behaviour to cover a larger or smaller search area.

Two experimental protocols were designed to meet these needs. In the first protocol (Table 1) a single parameter per experiment was modified, to understand the impact of each individual parameter on the final result.

The second protocol (Table 2) varied multiple parameters per experiment, in an attempt to understand more about the interaction of the parameters.

In each experiment, the simulation's experimental parameters were set, and each simulation was run for 1000 time steps. Experiments were run on the University of Reading's Condor High-Throughput Computing System. Starting positions for the flock members were selected from a uniform, random distribution to ensure uniform distribution of population at the start of each experiment. Each simulation was repeated ten times. The experiments using the parameters in Table 1 involved 100 birds in each experiment. This number of birds was selected as a suitable upper limit of robotic agents to be deployed in a real-world application.

The simulation allows multiple birds to occupy the same position, hence collision detection is not incorporated into the simulation. To avoid birds reflecting off the edge of the simulation, no bounding rules were incorporated. In comparison with the cornfield vectors of Heppner [13] and the Particle Swarm Optimisation method [3], the flock was not provided with a destination to move towards. As such, the birds were effectively instructed "to flock" continually, which is considered similar to the murmurations of starlings. Data was recorded on each bird's position and velocity per time step.

3.2. Metrics

From Parrish [14], there are two primary flocking behaviours used to identify the amount of flocking within an animal aggregation in nature, flock area and polarisation. "Flock area" is the area that fully contains all members of the flock. It is considered that a smaller flock area indicates a greater degree of flocking. Polarisation describes the alignment of the flock members. A group of highly polarised birds should demonstrate strong flocking behaviour.

However, there is no authoritative metric that can describe whether flocking is taking place. Currently the best approach involves human observation of a flock and qualitative decision making regarding whether the behaviour observed is flocking.

3.3. Metrics calculation

3.3.1. Area

At each time step in the simulation, a convex hull around the simulation participants is calculated, giving the area of the flock at each point.

Each set of experimental parameters was run ten times – "ten experiments" – this hence produced ten sets of results (areas). The area found at each time step across the simulations for that set of experimental parameters was then averaged. This average area is used in later analysis.

3.3.2. Polarisation

Flocking simulations configured with different sets of parameters were performed.

The data set for each simulation run contained the velocity information for each bird in the flock per time step. For each experiment, the standard deviation (SD) of the velocities of the birds in the individual flock is calculated, for each time step. This is to measure the spread of the individual bird velocities from the mean, with a closely aligned flock having a smaller SD.

For all of the ten runs, the SD of the individual flocks' SD is calculated. This measurement is then used in later analysis.

3.3.3. Commentary on approach

The standard deviation (SD) from the mean is widely used to indicate the degree of dispersion, and takes into account the deviation of every value from the mean [15]. With the polarisation measurement, we are interested in measuring a *spread of spreads*.

There are two phases to generating the final polarisation metric. The first part is termed the *intra-simulation phase*. In this first stage, the aim is to understand how well aligned the birds within an individual flock are at a single time step. This is done by measuring the SD of the individual flock at each time step throughout the course of one simulation run. In turn, this is repeated for each of the experimental parameters.

The second part to generating the polarisation metric is termed the *inter-simulation phase*. Simulations using the same experimental parameters are repeated ten times. This means that by the end of the intra-simulation phase there are ten results per time step, per set of experimental parameters. To combine these disparate results to create one final result per time step per set of experimental parameters, the SD of the intra-simulation results is calculated. Effectively, this is measuring the spread of the spread.

Table 1. Experimental parameters used for Protocol 1. Shading indicates parameters changed per experiment

	Experimental Protocol 1 Identifier														
Parameter	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Move Closer scaling factor	20	40	60	80	100	50	50	50	50	50	50	50	50	50	50
Move With scaling factor	50	50	50	50	50	20	40	60	80	100	50	50	50	50	50
Move Away scaling factor	25	25	25	25	25	25	25	25	25	25	10	20	30	40	50
Move Away Minimum Distance	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50
Maximum velocity	50	50	50	50	50	50	50	50	50	50	50	50	50	50	50
	Experimental Protocol 1 Identifier														
Parameter	16	17	18	19	20	21	22	23	24	25					
Move Closer scaling factor	50	50	50	50	50	50	50	50	50	50					
Move With scaling factor	50	50	50	50	50	50	50	50	50	50					
Move Away scaling factor	25	25	25	25	25	25	25	25	25	25					
Move Away Minimum Distance	20	40	60	80	100	50	50	50	50	50					
Maximum velocity	50	50	50	50	50	20	40	60	80	100					

Table 2. Experimental parameters used for Protocol 2

Experimental Protocol 2 Identifier	Number Birds	Move Closer scaling factor	Move With scaling factor	Move Away scaling factor	Move Away Minimum Distance	Maximum Velocity
1	50	100	25	5	20	10
2	25	80	50	10	40	30
3	5	60	75	15	60	50
4	100	40	100	20	80	70
5	75	20	25	25	100	90
6	50	5	50	30	20	110
7	25	100	75	35	40	10
8	5	80	100	40	60	30
9	100	60	25	5	80	50
10	75	40	50	10	100	70
11	50	20	75	15	20	90
12	50	20	75	15	20	110
13	5	100	25	25	60	10
14	100	80	50	30	80	30

4. Results & Discussion

The two key measurements presented are flock area and the standard deviation (SD) of flock polarisation. Results from the experimental protocols are presented in two sections. The results from the first experimental protocol (Table 1), where the experimental parameters are modified independently, are given in Section 4.1.

Section 4.2. contains the results from the second protocol (Table 2) where multiple experimental variables are changed together. A discus-

sion following integrating the results from the two protocols is given in Section 4.3.

Note that the area results presented throughout are negative. This is due to the direction used to calculate the convex hull.

4.1. Experiment Protocol 1: Vary individual experimental variables

4.1.1. Move Closer Scaling Factor

Figure 1a presents the results from changing the Move Closer scaling factor, s_c . As the Move Closer scaling factor value increases, flock area increases and flock polarisation remains consistent within a range. Tests for significance were performed between the flock polarisation results (t test, $P < 0.01$): differences between data sets were all significant except the results when $s_c = 20$ and $s_c = 60$. Figure 2 summarises the differences.

From equation 3, it can be seen that small values of the Move Closer scaling factor, s_c , results in larger values of the scaled distance, s_d . A larger scaled distance value results in birds being physically closer to one another following equation 4. This in turn results in a smaller flock area, as is observed as in Figure 1a.

As equation 3 results in a scaled distance value, it is expected that the polarisation of the flock will not be altered, as it is the position of the bird that is being modified. The flock polarisation remains consistent within a range, consistent with the observations.

It is considered that this would be a key variable to modify the area of, for example, a collection of autonomous robotic agents working together on a task. That is, if the task required covering different areas, for example in frontier exploration [16].

4.1.2. Move With Scaling Factor

In Figure 1b the results from modifying the Move With scaling factor, s_w , are presented. A large Move With scaling factor results in a less polarised flock, and minimally affects the flock area. Tests for significance were performed between the flock area results (t test, $P < 0.01$): dif-

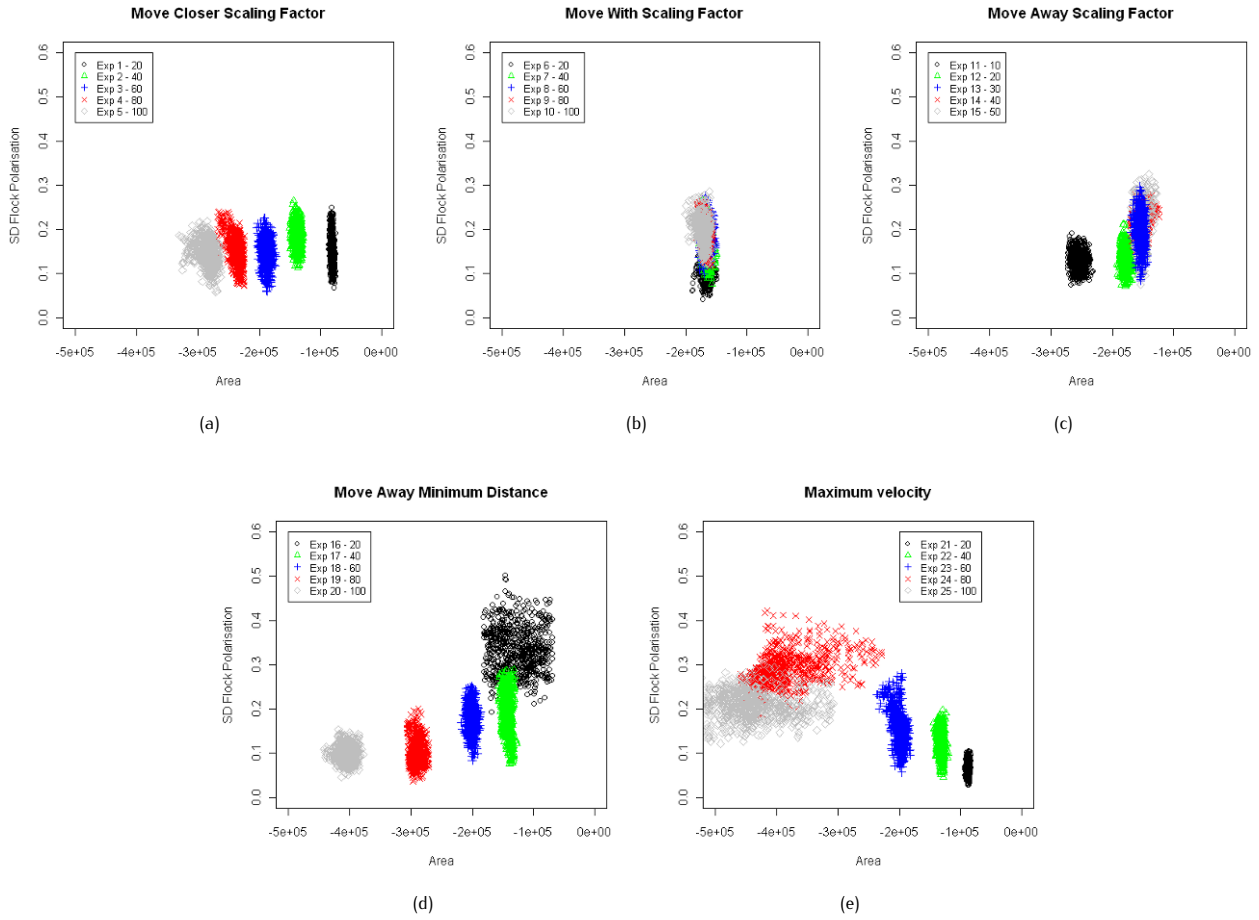


Figure 1. Scatter plots showing the flock area against polarisation results for time steps 249 – 999 (Experimental Protocol 1).

ferences between data sets were all significant except the results when $s_w = 20$ and $s_w = 40$ or $s_w = 60$. Figure 3 summarises the differences between results. Whilst this change in flock area is significant, it is small in comparison to other experiments.

The value scaled in 6 is the average flock velocity, \bar{v}_x . Equation 6 shows that a smaller Move With scaling factor, s_w , will result in a larger scaled average velocity, $s\bar{v}_x$. A larger scaled average velocity will have a larger effect on bird velocity, $b_{vx,t+1}$, following equation 7. A larger change to an individual bird's velocity, with said change being derived from the flock's average velocity, will have a greater effect on moving the bird, and cause the bird to move towards the flock's average velocity. Therefore the flock's polarisation should be greater when the Move With scaling factor is smaller, as is observed in Figure 1b.

Through the modification of this parameter, it should be possible to influence the polarisation behaviour of a swarm of autonomous robotic agents. There would be situations where having a highly polarised group of agents would be useful – for example, moving from one location to another in a collective fashion. Similarly, there would be other

situations where having the agents assume a less polarised “swarm” structure would be beneficial, such as in tasks involving foraging [17].

4.1.3. Move Away Scaling Factor

The observations from varying the Move Away scaling factor are shown in Figure 1c. It can be seen that as the scaling factor increases, both flock polarisation and flock area decreases.

The Move Away rule is intended to ensure that birds do not get too close to one another. This rule is only enforced when birds are within a defined distance from one another, as per 8. A limitation of the study is that it is unknown how many birds are impacted by the Move Away rule.

Equation 10 produces a scaled distance that is subtracted from a bird's velocity, to physically move the birds away from one another as required. If a smaller value of the Move Away scaling factor is used with this equation, the value of the scaled summed axis difference root, S_x , will be larger. When a larger scaled summed axis difference root value is used in 11, a smaller bird velocity will result. This will result to slow the

Move Closer Scaling Factor Polarisation results

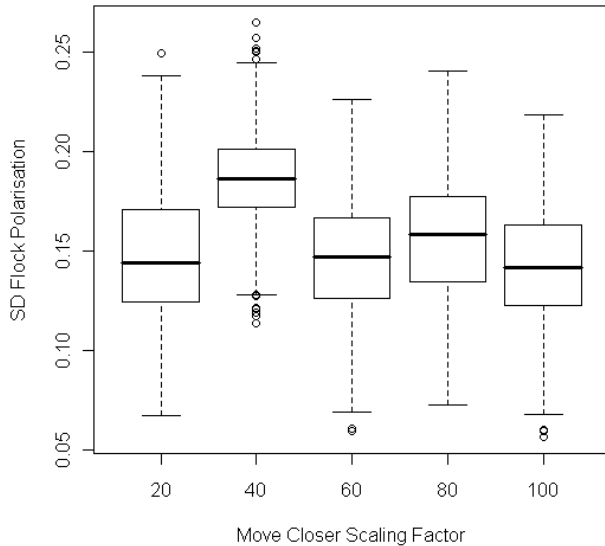


Figure 2. Polarisation results for the Move Closer scaling factor (Experimental Protocol 1)

Move With Scaling Factor Area results




Figure 3. Area results for the Move With scaling factor (Experimental Protocol 1)

bird and hence the bird will move a smaller distance. In turn, a smaller flock area will be observed, which is consistent with the results.

It is considered that the reduction in flock area then introduces positive feedback into the simulation. As flock area decreases as a result of the Move Away scaling factor increasing, it follows that the birds must be closer to one another. Due to the birds being in closer proximity to one another, more birds will be impacted by the minimum distance rule from 8. This means that the Move Away rule will then be applied to a greater proportion of the flock. As more birds have their velocity reduced by 11, the flock area will therefore decrease.

Further, the Move Away rule is performed after the Move With rule. The Move With rule helps to maintain polarisation within the flock. As a consequence, as more birds are impacted by the Move Away rule for the reasons explained, their velocities are modified, but the polarising effect of the Move With rule is not then reapplied to maintain flock polarisation. It is therefore considered that the flock's polarisation decreases as the Move Away scaling factor is increased, as can be seen in Figure 1c.

Specific settings of this parameter could prove useful in application areas where random search is required in a small area.

4.1.4. Move Away Minimum Distance

Figure 1d presents the observations from modifying the Move Away Minimum Distance value. When the Move Away Minimum Distance is increased, the area of the flock increases and flock polarisation increases also.

The Move Away Minimum Distance value affects the simulation in two ways. Firstly, the Move Away Minimum Distance value controls how many birds are affected by the Move Away rule, as per 8. However, the current study does not record how many birds are impacted by the Move Away rule at a given time. Therefore it is not possible to currently calculate precisely what impact the Move Away rule has on the flock as a whole.

However, if the flock area is large, it can be surmised that the average distance between birds will be large. Therefore if the Move Away Minimum Distance value is low, few birds will be impacted by the rule. On the contrary, if the Move Away Minimum Distance value is high and the flock area is low, it can be assumed that many flock members will be impacted by the rule.

The Move Away Minimum Distance secondly affects the simulation in equation 9, where the square root of the value is calculated. Larger values of Move Away Minimum Distance will result in larger values of the square root. In turn, larger values of the square root will result in larger values of the axis difference root, ad_{jx} . Use of a larger axis difference root value in 10 will result in a larger summed axis difference value, S_x . Carrying these considerations into equation 11, higher summed axis difference values will result in a smaller bird velocity at the next time step, $b_{vx,t+1}$.

Modifying this parameter may be relevant in situations where robotic agents need to "sweep" a specific area to achieve search goals, e.g. in searching a minefield for mines.

4.1.5. Maximum Velocity

In Figure 1e the observations from changing the Maximum Velocity of the flock are presented. A higher Maximum Velocity results in a larger flock area and less polarisation within the flock.

Birds with a higher Maximum Velocity are able to move a greater distance in a single time step. Therefore it is considered that this explains

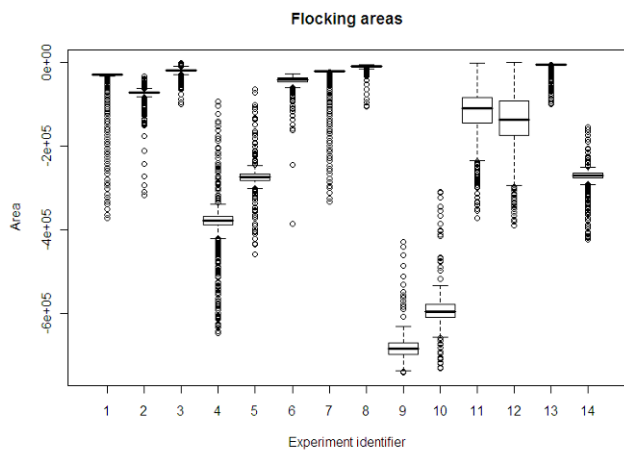


Figure 4. Results of Area experiment (Experimental Protocol 2)

why a larger flock area is observed when the birds have a higher maximum velocity: the birds are capable of moving further away from one another, resulting in larger observed flock areas.

Equations 13 and 14 are used to reduce a bird's velocity when it exceeds the experiment's maximum velocity v_{max} . However, when the Maximum Velocity is set to a high value, this means a bird can reach much higher velocities before its velocity is scaled. It is considered that higher values of v_{max} allow a bird to diverge more from the flock's average polarization. Therefore lower flock polarisation is expected with higher Maximum Velocity values, which is consistent with the observations.

In comparison with the Move Away Scaling rule, varying this parameter could be applicable when random search is required over a small area.

4.2. Experiment Protocol 2: Vary multiple experimental variables

Results for all experimental parameters are presented in Figure 4 for area, and Figure 5 for flock polarisation.

It is firstly considered that flock area is a more relevant measurement to use when smaller flocks are modelled. As commented previously, the model lacks any cornfield vector. The result of this is that the simulated flock is effectively directionless, and will continuously wheel around itself. If a cornfield vector were included, this would act as a destination for the flock, which would hence move in that direction. Further, the rules of the flocking model mean that the flock move towards the average central point of the population.

The consequence of this invalidates the velocity measurement for smaller flocks. As per the area metric description, a convex hull is used to identify the boundary of the flock. From this boundary, assuming that the birds are orientated towards the flock and that the birds are equally distributed throughout the flock, one can consider that these two factors will greatly reduce the flock polarisation. This could be further investigated through producing density maps, indicating how the birds are arranged within a flock.

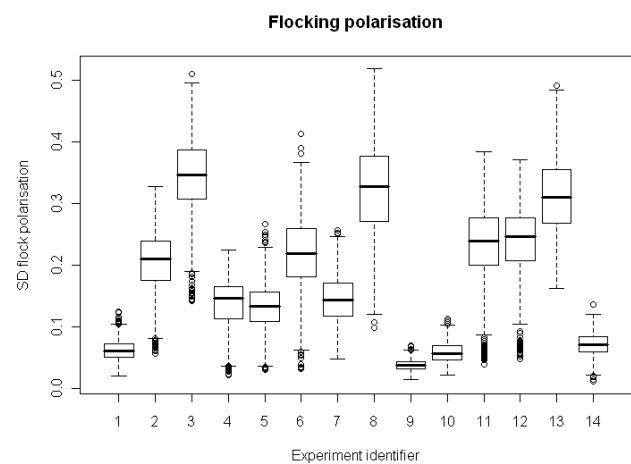


Figure 5. Results of Polarisation experiment (Experimental Protocol 2)

Figure 6 charts the velocity results of the flocks with the highest number of birds and those with the lowest. A comparison between the two flocks can be observed in Figure 6. Flocks with the highest number of birds (100) have the lowest standard deviation of polarisation, and hence are the most polarised. In comparison, the flocks with the lowest number of birds (5) are the least polarised.

This is a direct impact of the number of birds involved in the flock. When there are five birds in a flock, assuming that they are equally distributed and all orientated toward the flock centre, the angle of orientation of each bird should be greatly different from one another. From this, the variation from the mean angle of orientation will be high, and hence the polarisation will be low.

When there are larger number of birds in the flock, it is further considered that a "flock momentum" is established, based on the initial starting positions. That is: The initial starting positions of the flock are developed through a random-number generator. If the results of the random number generator were slightly non-uniform, this could adjust the centre of the flock. If 100 birds were heading to a point that was not the true centre of the flock, then this could develop a "momentum" in the flock, and through this, increase the polarisation of the flock members. This could be investigated further through comparing the distance travelled over time by the flock centre, to understand if there is an observable difference between flocks with large and small numbers of birds.

A scatter chart of average flock polarisation against the number of birds is presented in Figure 7. As shown, there is a good relationship between the number of birds in the flock and the polarisation observed. This supports the view that the number of birds in the flock is the primary influence on flock polarisation.

Figure 8 presents the results from the area experiments. Each chart is sorted by a different experimental parameter, with the minimum values on the left of the x axis, and maximum values on the right of the x axis. The first 250 time steps were omitted from the charts, to allow the simulated flocks to reach equilibrium.

In Figure 8a, the results from altering the Maximum Velocity parameter are shown. Through possessing a smaller Maximum Velocity parameter, a smaller final individual bird velocity will result. Therefore, one

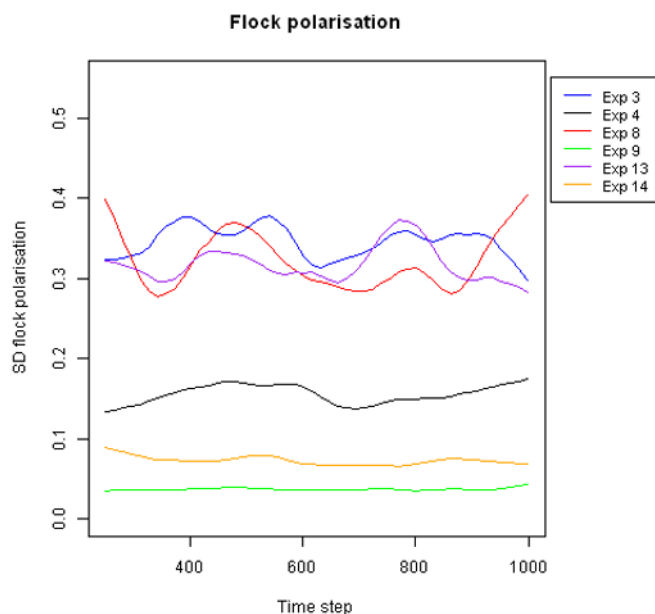


Figure 6. Fitted LOESS curves presenting the polarisation results for the experiments with the highest and lowest number of birds in the flock, experiments 3, 8 and 13, and experiments 4, 9 and 14 respectively. (Experimental Protocol 2)

expects that a smaller velocity will be produced, and hence a smaller final flock area will be observed. However, the experimental results shown in Figure 8a do not match this prediction, as a non-linear relationship is observed. On further investigation, when the results were compared against the number of birds in the simulation, it is found that the results follow flock size.

Figure 8b presents the impact of the number of birds on the flock area. As can be observed, the size of the flock is related to the average flock size.

The observations from modifying the Move Closer scaling factor are given in Figure 8c. A larger Move Closer scaling value should produce a smaller impact on the final velocity result. As such, having a larger Move Closer scaling value will only slightly affect the final velocity. In turn, the bird will have a larger final velocity. Hence it is predicted that this should lead to a larger final area: a larger final velocity should produce a larger flock area. However, the experimental observations do not match this prediction. Again, upon comparing the observed results with the number of birds in the simulation, it is found that flock size is the principle factor affecting behaviour.

Figure 8d presents the results from modifying the Move With scaling factor. Including a smaller Move With scaling factor should result in a greater reduction of the final velocity. This means the distance that birds could move away from one another in a single time step would be reduced. In turn, this would mean the final resulting flock area would be smaller. Hence a smaller Move With scaling factor should produce a smaller flock area, as confirmed by results generally shown in Figure 8d, with the exception of experimental identifier 4.

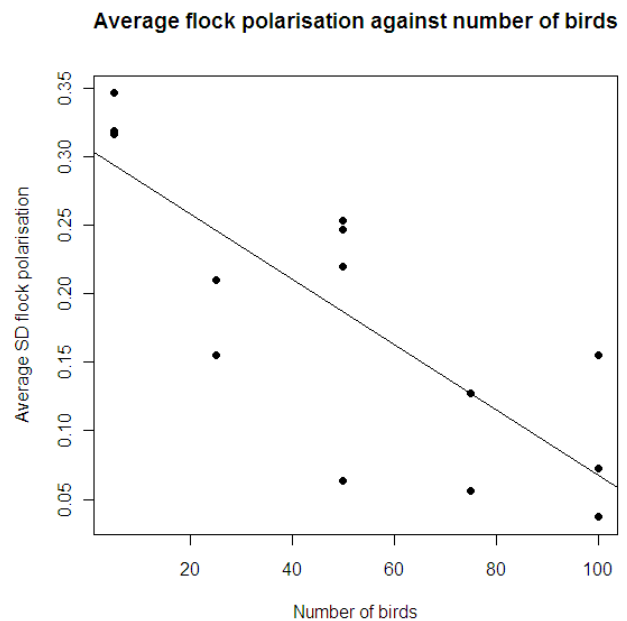


Figure 7. Scatter chart of flock polarisation against number of birds, with regression line. (Experimental Protocol 2)

In Figure 8e, the results of changing the Minimum Distance value are given. Larger values of Minimum Distance will mean that more birds would be impacted by the Move Away rule. Further, a large value of Minimum Distance leads to a small final velocity. Hence a smaller Minimum Distance value should result in a smaller flock area. However, this is in direct contrast to the observed simulated behaviour, as shown in Figure 8e. Upon closer inspection, it is found that the observed results follow the number of birds in the flock.

Figure 8f shows the results of using different Move Away scaling factor values. A larger value of the Move Away scaling factor produces a smaller final velocity. Birds possessing a smaller velocity will remain in a smaller total area, as they will not be able to move over a large distance in one time step. As a result, the final flock size will be smaller with a small Move Away value, as shown in Figure 8f.

4.3. Integration Experimental Protocols 1 and 2

Following the analysis of Experimental Protocols 1 and 2, it was questioned whether the results of Protocol 1, where the flocking simulation variables were independently varied, could be used to predict the observed behaviour in Protocol 2, where multiple flocking parameters were altered simultaneously.

Figure 9 presents the results from Experimental Protocol 2 where the number of birds used in the simulation is either $n = 50$ or $n = 100$. Note that the results for experiment 12 from Protocol 2 are omitted: the results for this experiment are considered very similar to those for Protocol 2, experiment 11, and did not add any additional value to this discussion.

The following steps were taken to perform the analysis:

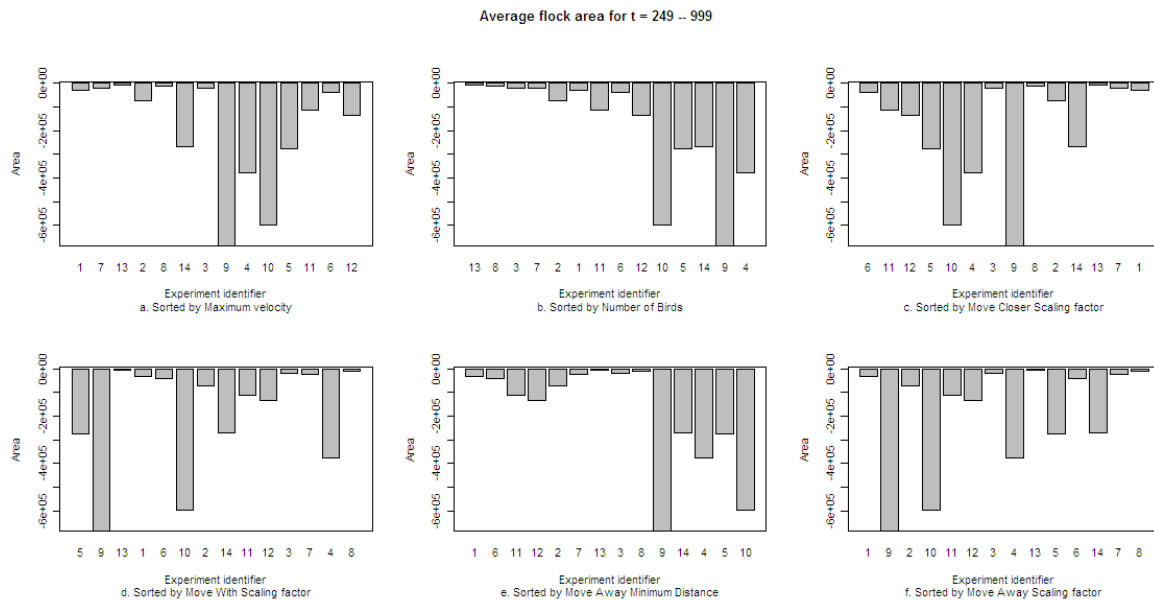


Figure 8. Bar charts showing the average flock area for time steps 249 – 999. Each chart is sorted by the stated experimental parameter, with the lowest value for each parameter on the left of the x axis, increasing in value to the right. (Experimental Protocol 2)

For each experiment in Experimental Protocol 2, the experimental parameters used were compared against the results from Protocol 1. Through this comparative step, an attempt was made to predict what flock behaviour would be observed when using the experimental parameters in Protocol 2.

As an example, the Maximum Velocity for experiment 14 in Experimental Protocol 2 was set to a value of $v_{max} = 30$. Table 2 shows that this is one of the lowest Maximum Velocity values used in Protocol 2.

This “low” Maximum Velocity value was then evaluated against the findings from Experimental Protocol 1 to understand: if a low Maximum Velocity value was used in an experiment as part of Protocol 2, what would be the predicted observed behaviour?

Figure 1e shows that a low Maximum Velocity value when varied independently in Experimental Protocol 1 results in a small flock area with high polarisation. It can be seen in Figure 9 that this is consistent with the observations from Protocol 2.

Through this approach, analysis can be made to understand the relative impact of the different experimental parameters on the observed behaviour. As an example, in experiment 1 of Experimental Protocol 2, the Move Away scaling factor is set to a value which would predict a large flock area. The other four experimental parameters in this experiment have values that predict a small flock area. When a small flock area is observed, it can therefore be considered that the combination of the four experimental parameters that predict a small flock area counteract the effect of the Move Away scaling factor. The approach can similarly be applied to the polarisation results.

In addition to the previous comment, it is considered that this analytical approach could have practical implications. In works such as [18, 19] where multiple robotic agents are used simultaneously to achieve a common goal, the ability to change the behaviour of a robotic flock is considered valuable. For example, the ability to modify the area a

robotic flock is examining in search-and-rescue type tasks could prove useful, depending on the size of the object being sought. Secondly the ability to change a group's dynamic from a polarised flock to a non-polarised swarm could be beneficial in multiple domains. These domains could include tasks where different agent dispersal methods are required [20, 21], or randomised information sharing, including distribution of tasks amongst agents [22].

Similar analysis from Experimental Protocol 2 experiment 9 indicates that the Move Away scaling factor and Move Away Minimum Distance parameters can counteract other values within a simulation to modify flock area. Further, Protocol 2 experiment 14 shows that values for the Move With scaling factor and Maximum Velocity can offset other experimental values to modify flock area. The results for the polarisation experiments are less clear: this would need to be investigated in greater depth in later studies to understand the impact on the observations.

5. Further Points For Discussion

There are a number of further points that require consideration ahead of applying the above findings to a domain. Firstly, two additional studies are considered of benefit to provide greater insight into the Reynolds flocking model described. Firstly, as highlighted previously, intermediate values from the simulation were not captured: for example, the number of birds affected by the Move Away rule was not recorded. In later studies, capturing this information would enable a better understanding of what proportion of the flock is being impacted by a particular rule.

Secondly, the Reynolds flocking parameter values selected for the experiments were chosen from a linear scale. Repeating the experiments

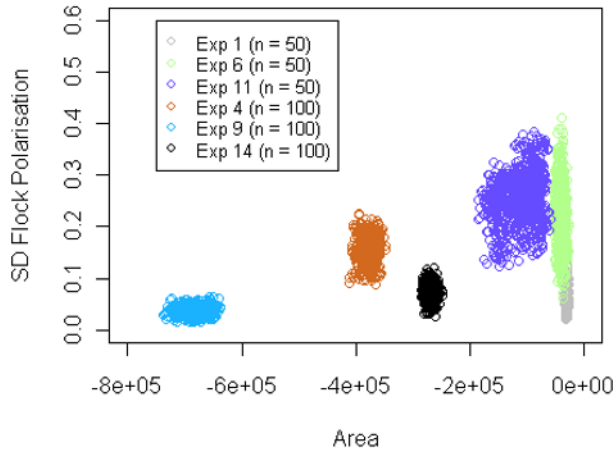
Experimental Protocol 2: $n = 50$ and $n = 100$ 

Figure 9. Results from Experimental Protocol 2 where the number of birds in the simulation either $n = 50$ or $n = 100$

with parameters selected from a logarithmic scale would provide additional insight, through exposing the flock to a greater range of parameter values.

The computational complexity of the calculations used was anticipated to be linear, with simulation times increasing proportional to the number of birds used. However, as presented in Figure 10, there is a polynomial relationship between the number of birds used and the time taken for the simulation to complete. Following on from the previous consideration, repeating the simulations but using an increased number of birds would provide further evidence with regards to the presence of a linear relationship. There may be a fixed simulation start-up cost, which once removed does reveal a linear relationship.

Further, it should be noted that the simulation operates in discrete steps. The simulated birds are stationary whilst they calculate their next positions through applying the described flocking rules. Once the positions have been calculated for the entire flock, the birds move to their new position simultaneously. Hence, this means that the birds are not continually moving. This would need to be investigated further ahead of applying the flocking behaviours to robotic agents that move along a continuous trajectory, as it would be important to understand if the flocking behaviours would need to be adapted for use by a robot that is already in motion. To this point, it is considered that implementing the above experiments in a robotics simulation package would provide additional understanding in this matter.

A number of simulation results in Experimental Protocol 2 do not agree with the predicted flock behaviour. Further work could investigate these discrepancies. However, on further inspection it is found that the number of birds in the simulation is the main factor affecting flock area. Hence one suggested approach would be to repeat the experimental procedures detailed, but keep the number of birds used in the simulation run constant. In conjunction with the findings from Experimental Protocol 1, this would allow a clearer understanding of the impact of combined flock parameters on flock behaviour.


Simulation run time

Figure 10. Simulation run times when experiments were conducted with varying numbers of birds. Data points are plotted with blue points; the black dashed line shows the results from applying a linear regression model; the red line shows the results from applying a polynomial regression model. R^2 values for the linear model are 0.94 and 1.00 for the polynomial model respectively.

6. Conclusion

An analysis of flocking simulations is presented.

Experimental Protocol 1 has described the impact the individual parameters within the flocking simulation have on the observed flock behaviour. The simulation results bear out the theory behind the flocking simulation in a decided manner.

From Experimental Protocol 2, it has been found that flock area is a suitable measurement for murmuring flocks, such as starling flocks. Further, it has been discussed that using flock polarisation with such flocks is an inadequate measurement. However, simulation results indicate that the number of flock members directly impacts the amount of polarisation observed.

One immediate extension to the above efforts is to extend the model to three dimensions. It is considered that there would not be a material impact to the above results, as the addition of a higher dimension should not greatly affect the operation of the rules and the interaction of the birds. However, it would be prudent to analyse these considerations, to ensure that this is the observed empirical result.

Studies have shown that birds maintain a network of around nine birds which they use to flock [23]. It would be beneficial to develop a new flocking model that uses these rules. This would allow a direct comparison between both the method of operation of the Reynolds model, as well as being able to compare the results from both models, and the behaviour observed from empirical flock studies.

In addition to the robotic applications of this research already mentioned, the findings of this research could be adapted to support investigations in other subjects. Examples could include physical and cognitive herding, human behaviour following media influences, and

brand purchasing behaviour post exposure to different advertising techniques.

References

- [1] E. Shaw, In: L.R. Aronson (Ed.), *Development and Evolution of Behavior* (W.H. Freeman and Company, San Francisco, 1970) 452
- [2] C.W. Reynolds, *SIGGRAPH Comput. Graph.* 21, 25 (1987)
- [3] J. Kennedy, R. Eberhart, *Proc. 1995 IEEE Int. Conf. Neural Networks*, 1942 (1995)
- [4] A. Schadschneider, W. Knospe, L. Santen, M. Schreckenberg, *Physica A*, 346, 165 (2005)
- [5] S. Thrun et al, *STAR Springer Tracts in Advanced Robotics Series* 36, 1 (2007)
- [6] C. Stiller, G. Farber, S. Kammel, 2007 *IEEE Intelligent Vehicles Symposium*, 215 (2007)
- [7] E. Guizzo, *IEEE Spectrum*, 45, 26 (2008)
- [8] S. Barbera et al., 12th *Int. Conf. Computer Modelling and Simulation (ICCMS)*, 104 (2010)
- [9] R.A. Russell, D. Thiel, R. Deveza, A. Mackay-Sim, *Proc. 1995 IEEE Int. Conf. Robot. Automat.*, 556 (1995)
- [10] H. Ishida, H. Tanaka, H. Taniguchi, T. Moriizumi, *Auton. Robot.*, 20, 231 (2006)
- [11] W. Baetz, A. Kroll, G. Bonow, *Proc. 2009 IEEE Int. Conf. Robot. Automat.*, 2773 (2009)
- [12] Q. Stephen-Hassard, E. Grabbe, *IEEE Oceans*, 735 (1977)
- [13] F.H. Heppner, U. Grenander, In: S. Krasner (Ed.), *The Ubiquity of Chaos* (AAAS Publications, Washington, 1990)
- [14] J.K. Parrish, W.M. Hammer, C.T. Prewitt, *Animal Groups in Three Dimensions* (Cambridge University Press, United Kingdom, 1997)
- [15] K.A. Stroud. *Engineering Mathematics*, 5th edition (Palgrave Macmillan, New York, 2001)
- [16] Y. Wang, A. Liang, H. Guan, 2011 *IEEE Symposium on Swarm Intelligence (SIS)*, 1 (2011)
- [17] A. Campo et al., *Biological Cybernetics*, 103, 339 (2010)
- [18] R. O'Grady et al., In: G. Kampis, I. Karsai, István, E. Szathmáry (Ed.), *Lecture Notes in Computer Science: Advances in Artificial Life. Darwin Meets von Neumann* (Springer, Berlin / Heidelberg, 2011) 165
- [19] K. Muecke, B. Powell, In: K. Mehrotra, C. Mohan, J. Oh, P. Varshney, M. Ali (Ed.), *Modern Approaches in Applied Intelligence* (Springer, Berlin / Heidelberg, 2011) 266
- [20] J. Pugh, A. Martinoli, *Swarm Intelligence*, 3, 203 (2009)
- [21] M.A. Montes et al., In: M. Dorigo et al. (Ed.), *Lecture Notes in Computer Science: Swarm Intelligence* (Springer, Berlin / Heidelberg, 2010) 251
- [22] F. Ducatelle, G. Di Caro, C. Pinciroli, L. Gambardella, *Swarm Intelligence*, 5, 73 (2011)
- [23] A. Cavagna et al., *PNAS*, 107, 11865 (2010)