Synthetic Flocking and Bioinspiration

B. D. Flom

***Abstract*—There has been a growing interest in understanding collective behavior in nature, in particular, how local interactions between animals give rise to seemingly coordinated behavior for an entire group. We study position data collected every 20 milliseconds from a flock of ten trained pigeons, fitted with GPS-loggers (Nagy et. al., Nature, 2010). We hypothesized that the flock motion can be modeled as a problem of networked interaction with pigeons engaged in a particular pursuit strategy with respect to others in the flock. We tested this hypothesis and found evidence suggesting that some pigeons implement a constant bearing pursuit strategy.**

I. Introduction

A. Background

A growingarea of research is that of multi-robot systems, with applications ranging in scale from microbots that could potentially search under rubble after an earthquake for survivors [1], to unmanned air vehicle swarms for search and rescue, reconnaissance, and defense [2].

One major challenge in the design of multi-robot systems is motion planning, since concurrently planning actuator inputs for each robot causes the computational complexity of the planning task to increase exponentially with the number of robots [3].

As a result, many researchers are studying multi-robot motion planning from the perspective of distributed control, where each robot implements its own feedback motion control law, based on limited state information of other robots in the group.

From an engineering standpoint, the task of designing separate motion planning algorithms for the individual robots such that an entire group is able to execute a desired coherent motion may seem daunting. However, we observe many instances in nature where groups of animals are able to maintain sophisticated collective behavior, despite the absence of an external decision-maker relaying motion commands to every animal in the group. Examples range from flocks of birds or schools of fish maintaining a seemingly coordinated formation [4], to the remarkably sophisticated ants that construct colonies and networks of highways [5].

We ultimately seek to understand how local interactions between agents, be they animals or robots, can give rise to seemingly coordinated behavior for an entire group. To do so, we must understand what underlying mechanism drives these local interactions.

One plausible local interaction is called constant bearing (CB) pursuit, where one agent, designated the *pursuer*, seeks to maintain a constant angle between its direction of motion, or heading, and the bearing towards another agent, designated the *evader*. Figure 1 illustrates CB pursuit, where $\boldsymbol{r}_p$ is the position of the pursuer, $\boldsymbol{r}_e$ is the position of the evader, $\boldsymbol{x}_p$ is the instantaneous direction of motion of the pursuer, and α is the constant bearing angle.

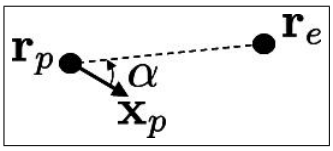

Fig. 1. CB pursuit (figure from [8])

Previous research has shown that CB pursuit, which some have suggested is the strategy used by falcons to pursue prey [7], can lead to rich behaviors for an entire group of agents all executing CB pursuit simultaneously [9], [10], [11]. For example, figure 2 illustrates the trajectories of four particles, indexed $1, 2, 3, 4$, all implementing CB pursuit such that particle i is pursuing particle $i+1$ modulo 4, i.e. particle 4 pursues particle 1 (this distribution of pursuit strategies is called *cyclic constant bearing pursuit*).The four plots correspond to different choices in bearing angles, and exhibit rectilinear, circling, expanding, and spiraling behaviors.

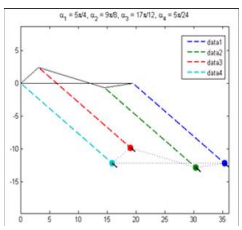

(a) Rectilinear

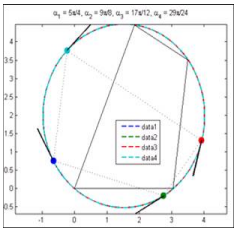

(b) Circling

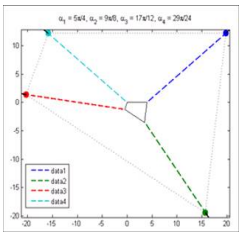

(c) Expanding

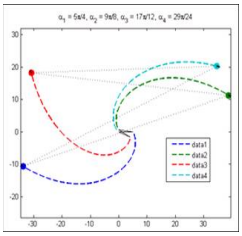

(d) Spiraling

Fig. 2. Examples of the variety of collective motion patterns induced by cyclic constant bearing pursuit, a distributed control algorithm based on local interactions between agents (figures from [8])

Observations of the collective in figure 2 give the impression there is a prescribed behavior for the entire group. However, each particle uses a feedback control law which only uses state information about one other particle.

Manuscript received August 1, 2011. This material is based upon work supported by the National Science Foundation under Grant No. 1063035.

B. D. Flom is with the Department of Electrical and Computer Engineering, University of Maryland, College Park, MD, 20742 USA e-mail: beflom@gmail.com.

B. Project Summary

In this work, we study position data collected by Nagy et al. every 20 milliseconds from a flock of ten trained pigeons, fitted with GPS-loggers, reported in [6]. In an analysis conducted by Nagy et al., it was shown that some pairs of pigeons had strongly correlated delayed velocity directions. That is, a pattern emerged where one pigeon's heading lagged behind another pigeon's heading with a fairly consistent delay.

Remark 1. The correlations in fact formed a hierarchy, as shown in figure 3. The pigeons are tagged with a letter, and the numerical label on an arrow pointing from Pigeon X to Pigeon Y is the delay corresponding to the peak correlation in heading, with Pigeon Y lagging behind Pigeon X.

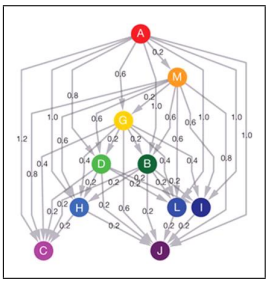

Fig. 3. Nagy leader-follower hierarchy for the pigeon flock (figure from [6]). See Remark 1 for an explanation of the network.

Nagy et al. proposed that this hierarchy of correlated velocities was the result of a more fundamental leader-follower hierarchy, so that the higher up in figure 3 a pigeon was, the more of a leader it was. In particular, pigeon A was hypothesized to be the leader of the entire flock since no pigeons are above it.

There are two interesting, yet puzzling observations to be made about figure 3. First, the delay times are not transitive. For example, pigeon M's heading has a peak correlation with pigeon A's heading delayed by 0.2 seconds, and pigeon G's heading has a peak correlation with pigeon M's heading delayed by 0.2 seconds, yet pigeon G's heading has a peak correlation with pigeon A's heading delay by 0.6 seconds, when one might expect it to be 0.4 seconds.

Another interesting feature of the network is that pigeons farther down the hierarchy have a tendency to have their velocities correlated with more "leader" pigeons than pigeons higher up the hierarchy. How many leaders can a single pigeon follow concurrently?

From these observations, it is not clear whether aligning directions of motion is the cause of the collective flock behavior, or rather, is a consequence of some more fundamental mechanism at work. For example, suppose the pigeon flock were to exhibit motions as in figure 2(a). If one were to study the position data, they might suggest that the pigeons were trying to align their headings to maintain a collective rectilinear motion. However, as was indicated earlier, the underlying mechanism is actually cyclic CB pursuit.

From the intuitive nature of CB pursuit, and previous research showing the richness of collective behavior driven by local CB pursuit interactions, we hypothesized that CB pursuit governs the motion of the pigeon flock. We also hypothesized that the leader-follower hierarchy proposed by Nagy et al. is consistent with CB pursuit, so that a follower pigeon would be following a leader with a defined offset angle.

To test these hypotheses, it was necessary to estimate velocities from position data that was collected by Nagy et al. We attempted to use an optimization-based approach to a fit the curve to the data that minimized a cost function which we discuss later. This curve would approximate the trajectory a pigeon actually flew. The optimization approach proved unsuccessful and will be discussed later in this paper.

Consequently, we based our CB pursuit analysis on the velocity data reported by Nagy et al., which they obtained by first averaging the position data and then using cubic B-spline interpolation. We found that, in general, the hierarchy proposed by Nagy et al. was not strongly consistent with CB pursuit. On the other hand, for certain pairs of pigeons, not necessarily a part of the hierarchy proposed by Nagy et al., the data was consistent with CB pursuit.

C. Paper Outline

In Section II, we discuss the theory behind the optimization-based curve estimation algorithm we attempted to use on the pigeon data. In Section III, we state the results of the curve estimation algorithm. In Section IV, we report the results of the CB pursuit test on the pigeon data. In Section V, we propose possible theories that could explain the lack of consistency between CB pursuit and the pigeon data, in particular, the hierarchy proposed by Nagy et al. In Section VI, we conclude with some remarks about future work which can test the theories proposed in Section V.

II. Curve Estimation

A. Natural Frames

In this section, we review some terminology and notation in the differential geometry of curves. The purpose of the following exposition is to be able to precisely describe the objects we study in the remainder of the section, that is, curves we are seeking to fit to a given sampled data set.

Consider a twice continuously differentiable (C^2) curve $\Gamma \subset \mathbb{R}^3$, with parameterization $t \mapsto \boldsymbol{\gamma}(t) \in \Gamma,\ t \in I \subset \mathbb{R}$, where I is an open interval. For ease of exposition, we refer to the parameter t as time. Let $t \mapsto \boldsymbol{T}(t) \in \mathbb{R}^3$ be the unit tangent vector to $\boldsymbol{\gamma}(t)$, and let $t \mapsto \nu(t) := ||\dot{\boldsymbol{\gamma}}(t)|| \in \mathbb{R}$, so that

$$\dot{\boldsymbol{\gamma}}(t) = \nu(t)\boldsymbol{T}(t). \tag{1}$$

For a particle traveling along the curve Γ at a timescale corresponding to $\boldsymbol{\gamma}(t)$, we can think of $\nu(t)$ as the speed of the particle at time t, so henceforth, we refer to ν as the speed.

A *continuously differentiable vector field on the curve* Γ is a continuously differentiable vector valued mapping $\Gamma \ni \boldsymbol{p} \mapsto \boldsymbol{v} \in \mathbb{R}^3$. An *orthonormal frame on the curve* Γ is a set of three continuously differentiable vector fields $\{\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3\}$ on Γ, with the property that for all $\boldsymbol{p} \in \Gamma,\ ||\boldsymbol{v}_i(\boldsymbol{p})|| = 1$ for $i \in \{1, 2.3\}$, and $\boldsymbol{v}_i \cdot \boldsymbol{v}_j = 0$ for $i, j \in \{1, 2, 3\}$ with $i \neq j$.

Henceforth, a "frame" is understood to mean an "orthonormal frame."

For a given parameterization of Γ, say $\boldsymbol{\gamma}(t)$, we can consider how a given frame *evolves* along the curve with respect to the parameterization. By a small abuse of notation, we define $\boldsymbol{v}_i(t)$ to be the composition of $\boldsymbol{v}_i$ with the parameterization $\boldsymbol{\gamma}(t)$. Then, the frame evolves with the dynamics:

$$\begin{aligned} \dot{\boldsymbol{v}}_1(t) &= \omega_z(t)\boldsymbol{v}_2(t) - \omega_y(t)\boldsymbol{v}_3(t) \\ \dot{\boldsymbol{v}}_2(t) &= -\omega_z(t)\boldsymbol{v}_1(t) + \omega_x(t)\boldsymbol{v}_3(t) \\ \dot{\boldsymbol{v}}_3(t) &= +\omega_y(t)\boldsymbol{v}_1(t) - \omega_x(t)\boldsymbol{v}_2(t) \end{aligned} \tag{2}$$

The vector $\boldsymbol{\omega}(t) := [\omega_x(t)\ \omega_y(t)\ \omega_z(t)]^T \in \mathbb{R}^3$ is commonly referred to as the *angular velocity of the frame at time t*.

Observe that in the system (2), the evolution of *each* vector field depends on the state of the other two vector fields. It can be shown [12] that there exists a choice of frame $\{\boldsymbol{T}, \boldsymbol{M}_1, \boldsymbol{M}_2\}$, with the vector field $\boldsymbol{T}$ being the field of unit tangent vectors along the curve, such that the frame evolves with the dynamics:

$$\begin{aligned} \dot{\boldsymbol{T}}(t) &= \nu(t)(k_1(t)\boldsymbol{M}_1(t) + k_2(t)\boldsymbol{M}_2(t)) \\ \dot{\boldsymbol{M}}_1(t) &= -\nu(t)k_1(t)\boldsymbol{T}(t) \\ \dot{\boldsymbol{M}}_2(t) &= -\nu(t)k_2(t)\boldsymbol{T}(t), \end{aligned} \tag{3}$$

which is simpler than the more general system (2). Such a frame is called a *natural frame*, and the functions $t \mapsto k_1(t), k_2(t) \in \mathbb{R}$ are called the *natural curvatures at time t*.

Remark 2. Conversely, if we do not know the frame values at all times *a priori*, but rather have ν, k_1, and k_2 prescribed for all t, then given an inital frame at time t_0, $\{\boldsymbol{T}(t_0), \boldsymbol{M}_1(t_0), \boldsymbol{M}_2(t_0)\}$, the evolution of the frame is then the unique solution to the system (3). This point is important for the next section.

B. Optimal Curve

We wish to find the curve that "best" fits a given data set consisting of positions in three dimensions. We shall discuss in the next section what is precisely meant by best, but for now suppose there is a cost function J associated with each candidate curve so that we seek the curve which minimizes J.

The space of all C^2 curves is infinite-dimensional. In our approach, we discretize the problem by approximating $\boldsymbol{\gamma}$ as having uniform piecewise-constant curvatures and speeds over the same time intervals. In other words, we are partitioning the interval I with the points $t_1, \cdots, t_n$, and assuming that ν, k_1, and k_2 are constant over every interval $[t_i, t_{i+1})$.

Then, the total number of parameters which uniquely characterizes a curve are the $3(n-1)$ constant curvatures and speeds, as well as the 3 scalar initial conditions, corresponding to the vector initial condition $\boldsymbol{\gamma}(t_0)$, and the 3 scalar initial conditions which can uniquely characterize the orthonormal frame $\{\boldsymbol{T}(t_0), \boldsymbol{M}_1(t_0), \boldsymbol{M}_2(t_0)\}$.

Our cost J, to be defined shortly, is a scalar function of the $(3n+3)$ constant curvature and speeds, and initial conditions. To minimize it, we make use of the MATLAB optimization toolbox function *fminunc*, which is an unconstrained optimization solver based on the steepest descent method. In principle, the function should search over the space of curvatures, speeds, and initial conditions, (which, by Remark 2, uniquely characterize the curves), and stop when it attains the minimum cost, which corresponds to the optimal curve.

C. Cost of Estimated Curve

We need a metric for measuring how "good" a candidate curve fits the data. The simplest metric would be to take the sum of squares of the errors between the candidate curve and the data points. However, because the data has noise, we actually do not want the curve that minimizes the fit error, which could look something like figure 4(a). Instead, we a want a curve that looks more like that in figure 4(b).

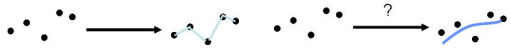

(a) Sum of squares of errors (b) Alternative

Fig. 4. Plausible metrics for measuring how "good" a candidate curve fits a data set

Thus, our cost function is a sum of the fit error and a term that penalizes the net change in speed and curvatures:

$$J = \sum_{j=1}^{n} ||\gamma(t_j) - \gamma_j||^2 + \lambda \int_0^{t_n} \left(\dot{k}_1^2(\sigma) + \dot{k}_2^2(\sigma) + \dot{\nu}^2(\sigma)\right) d\sigma, \tag{4}$$

where n is the number of data points, γ_j is the j'th data point, $\gamma(t_j)$ is the point on the estimated curve corresponding to γ_j, and λ weights the penalty discussed earlier. We need to specify λ.

Remark 3. When $\lambda = 0$, the cost function J reduces to the sum of squares of the errors. We would expect this choice of λ to result in the minimal curve having the most frequent sharp turns. On the other hand, it is not obvious what the minimal curve would look like for λ arbitrarily large. Indeed, some possibilities include a straight line, a circle, and piecewise arcs of constant curvature and constant speed.

To find an appropriate λ, we use a technique called *cross validation* (CV), which can be described as follows. We want to see which choice of λ, out of say m possible values, is the correct λ to use for the cost function.

Suppose we have n data points. Using each of the m choices for λ, we obtain m well-defined cost functions from equation (4). Then, we use the curve estimation algorithm discussed in Section II-B on $(n-1)$ data points, and measure the error between the estimated curve and the excluded data point. For each choice of λ, we repeat this process for each of the n points and sum the errors for each excluded point. Then, the λ we select is that which minimizes the total error in estimating the excluded points.

We call the total error in one step of cross validation the *validation cost* corresponding to a particular choice of λ.

III. Results from Curve Estimation

A. Cross Validation

The results reported by Nagy et al. were based on approximately 33 minutes of flight. Due to the computational demand of the cross validation procedure for the curve estimation algorithm, we restricted our analysis to a 77 second duration of flight. We ran the cross validation procedure, as discussed in the previous section, for pigeons A and M for this flight duration.

The initial set of candidate $\lambda's$ was $\{e^n \mid n \in \{-14, -13, \cdots, 1\}\}$. This scale was chosen because from experience, we found that the cost corresponding to small values of λ is more sensitive to perturbations than for large λ. Figure 7 shows, for pigeon M, the validation cost for each of the mentioned candidate values for λ. The λ corresponding to the smallest validation cost was e^{-4}.

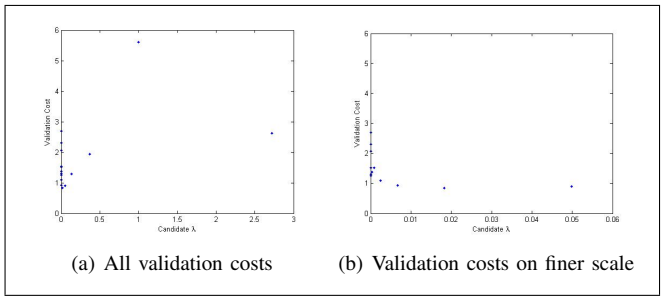

(a) All validation costs (b) Validation costs on finer scale

Fig. 5. Validation cost as a function of λ for pigeon M

Figure 6 shows, for pigeon A, the validation cost for each of the mentioned candidate values for λ. The λ corresponding to the smallest validation cost was e^{-8}. However, observe that the order of magnitude for the minimum validation costs for pigeon A is several times larger than the minimum validation cost for pigeon M. We hypothesized that this may be the result of the curve estimation algorithm not converging to the global minimal curve for pigeon M.

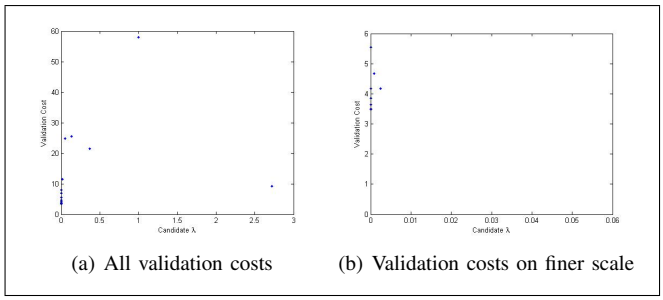

(a) All validation costs (b) Validation costs on finer scale

Fig. 6. Validation cost as a function of λ for pigeon A

To test this hypothesis, we examined the minimal curves at steps in the cross validation. Figure 7(a) shows the minimal curve (corresponding to the optimal λ) fit through pigeon A's data in the 20'th step of the cross validation (i.e. the 20'th point is removed). Intuitively, it is obvious that this curve is not a good fit. Contrast it with the minimal curve fit through pigeon M's data, illustrated in figure 8, which seems like a much more reasonable fit.

Since the optimal λ for pigeon A, denoted λ_a, is orders of magnitude smaller than the optimal λ for pigeon M, denoted λ_m, it seemed worthwhile to examine a plot of a minimal curve for pigeon A data using λ_m. Figure 7(b) shows just that, and indeed, the curve is still a bad fit, indicating a deeper problem in the curve estimation algorithm with pigeon A's data.

Therefore, we conducted our CB pursuit analysis based on the velocity data reported by Nagy et al., obtained by averaging and interpolation methods.

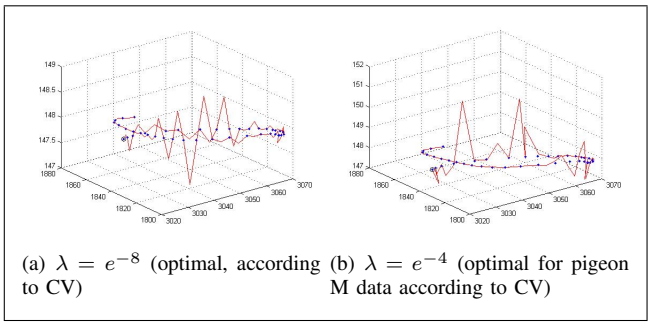

(a) $\lambda = e^{-8}$ (optimal, according to CV) (b) $\lambda = e^{-4}$ (optimal for pigeon M data according to CV)

Fig. 7. Estimated curves fit to pigeon A data corresponding to two choices of λ

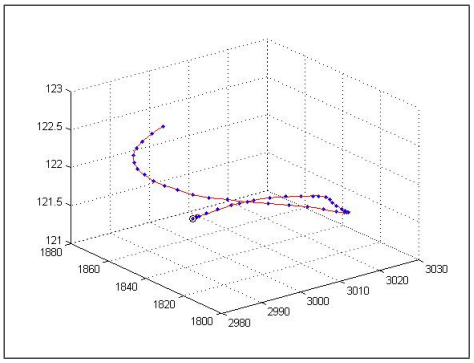

Fig. 8. Estimated curve fit to pigeon M data corresponding to $\lambda = e^{-4}$ (optimal according to CV)

IV. Preliminary Results of CB Pursuit Analysis

Using the velocity calculations reported by Nagy et al., we computed, at each time step, the bearing angle between a candidate follower's velocity vector and the directional vector pointing from the candidate follower to the candidate leader (i.e. using the notation of figure 1, we computed $\alpha = \cos^{-1}(\boldsymbol{x}_p \cdot (\boldsymbol{r}_e - \boldsymbol{r}_p))$.

We then plotted histograms of the distribution of bearing angles for different pairs of pigeons. For instance, in order to test whether Pigeon X was using CB pursuit to follow Pigeon Y, we set $\boldsymbol{r}_e$ to be the position of Y and $\boldsymbol{r}_p$ to be the position of X. If X were indeed using CB pursuit to follow Y, then we would expect a peak at some bearing angle in the histogram, which would indicate that the offset angle is frequently near the desired bearing angle.

It is not *a priori* clear how to precisely define a peak in the histogram. Rather, at this preliminary stage, we are making qualitative observations and using intuition. As we uncover

more information, we shall revise our methodology to make it more quantitative and precise.

According to the leader-follower hiearchy proposed by Nagy et al., the seniormost leader in the flock is pigeon A, followed by pigeon M. Figure 9 shows a histogram of bearing angles under the assumption that M is pursuing A. The bearing range is from 0 to π radians, and each bin has size $\pi/18$.

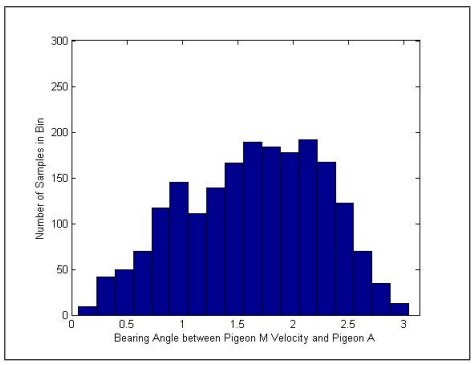

Fig. 9. Distribution of bearing angles under the hypothesis that M pursues A

While one cannot necessarily rule out the possibility that M is following A using CB pursuit, assuming that M is pursuing G or I gives rise to a better defined peak, making those cases more consistent with CB pursuit, as can be seen in figure 10. Note, however, that this is inconsistent with the hierarchy proposed by Nagy et al.

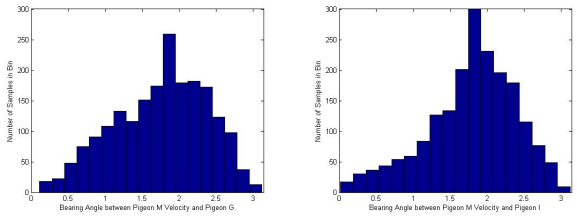

(a) M pursues G (b) M pursues I

Fig. 10. Distribution of bearing angles under the hypothesis that M is a follower

Conversely, if we assume that either G or I pursues M, we also see a peak in the bearing angle distribution, as shown in figure 11. This result is interesting because in general, CB pursuit is not commutative, i.e., if X pursues Y using CB pursuit, in general nothing can be concluded as to whether Y pursues X.

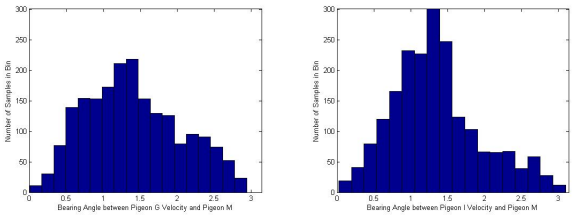

(a) G pursues M (b) I pursues M

Fig. 11. Distribution of bearing angles under the hypothesis that M is a leader

V. Discussion

A. Curve Estimation

The optimization-based curve estimation algorithm we used did not converge properly for pigeon A. The optimization algorithm uses a steepest-descent approach, so it is likely that it converged to a local minimum of the cost function (4) due to the cost not being convex. There are techniques for improving the performance of nonconvex optimization algorithms, often times trying many random initial conditions. These approaches are computationally demanding, and we lacked the resources to do this effectively.

B. CB Pursuit Analysis

We hypothesized that the underlying mechanism governing local interactions between pigeons in a flock is CB pursuit, and moreover, that the leader-follower hierarchy proposed by Nagy et al. is consistent with CB pursuit.

The distribution of bearing angles for certain pairs of pigeons, such as M pursuing A, does not show a clear consistency with CB pursuit. However, since at this stage we do not yet have a rigorous metric that measures likelihood of CB pursuit, we are not yet ready to dismiss our hypothesis.

On the other hand, for some pairs of pigeons not related to the Nagy hierarchy, our results are fairly consistent with CB pursuit.

One possible explanation for this phenomena is that follower pigeons might not be following a unique leader pigeon for the duration of the entire flight. For example, suppose pigeon A was indeed the leader of the entire flock, and suppose pigeon H is one step higher than pigeon C on the hierarchy. If A happened to be the pigeon nearest C, why would C attempt to still follow H if it knows that A is the seniormost leader? It seems more intuitive that it would be sufficient for a pigeon to follow any pigeon higher than it on the hierarchy, so that eventually the network leads to the leader A.

Another interesting result in our analysis was an apparent commutativity between certain pairs of pigeons, in particular, the pairs (M,G) and (M,I). Again, there is no reason to, in general, expect that if X follows Y using CB pursuit, that there should be any numerical suggestion that Y follows X using CB pursuit. An explanation for this phenomena is that perhaps, certain pairs of pigeons take turns leading versus following. This might be motivated by a desire to allow pigeons to draft behind others, using a strategy similar to bicyclists in a competition.

VI. Conclusion

A. Summary

The ultimate objective of this research program is to understand how local interactions between agents, be they animals or robots, can give rise to seemingly coordinated behavior for an entire group. Such an approach to multi-robot systems would facilitate the development of motion planning algorithms potentially more effective than concurrently executing motion planning tasks on an external system.

In the present work, we studied a particular example of collective motion, that of pigeon flocks. In particular, we tested the hypothesis that a local interaction known as constant bearing pursuit is the mechanism governing the behavior of the entire flock.

There were two aspects to this project. The first was to reconstruct the trajectory a pigeon traversed, from measured position data, and use the parameterization of the curve to compute the pigeon's velocity at different time steps. The reconstruction of the trajectory was attempted by using an optimization-based algorithm, that frequently converged to local minima due to the objective function being nonconvex. Consequently, we used the velocity data computed by Nagy et al., derived from an averaging and interpolation algorithm.

The second aspect of this project was to use the velocity values to compute the distribution of bearing angles corresponding to a hypothesized follower pigeon pursuing a hypothesized leader pigeon. The distribution was used to analyze whether constant bearing pursuit might be a strategy pigeons use in a flock, and furthermore, how consistent a pursuit hierarchy was with the leader-follower hierarchy proposed by Nagy et al.

Our preliminary results suggest that select pairs of pigeons support the hypothesis of CB pursuit, but there is not enough evidence to suggest that the Nagy hierarchy is consistent with CB pursuit. More complete results might allow an alternative hierarchy, where followers do not necessarily follow a unique pigeon for an entire flight.

B. Future Work

The optimization-based curve estimation algorithm must be modified to improve performance. In particular, we must develop techniques to prevent the algorithm from converging to a local minimum. That the curve estimation algorithm requires shorter flight durations is not an inconvenience, as we must examine short durations of flight to test the hypothesis that a single follower might switch between leaders, as discussed in the previous section. We will also look for evidence of certain pairs of pigeons switching between leader and follower roles. Finally, we will consider alternative pursuit strategies as the mechanism governing the behavior of the pigeon flock.

Acknowledgment

The author would like to thank his graduate student mentors B. Dey, K. S. Galloway, and M. Mischiati, as well as his adviser Professor P. S. Krishnaprasasad for their extremely helpful guidance and mentoring throughout this project.

References

[1] D. Pescovitz, "Robot Bugs: Handy, Not Icky," *Popular Science*, published 8 May 2004, accessed 1 Aug. 2011. <http://www.popsci.com/scitech/article/2004-05/robot-bugs-handy-not-icky>

[2] R. W. Beard, T. W. McLain, D. B. Nelson, D. Kingston and D. Johanson, "Decentralized cooperative aerial surveillance using fixed-wing miniature UAVs," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1306-1324, 2006.

[3] J. P. Desai, J. P. Ostrowski and V. Kumar, "Modeling and Control of Formations of Nonholonomic Mobile Robots," *IEEE Trans. on Robotics and Automation*, vol. 17, no. 6, pp. 905-908, 2001.

[4] J. Toner and Y. Tu, "Flocks, herds, and schools: A quantitative theory of flocking," *Physical Review E*, vol. 58, no. 4, pp. 4828-4858, 1998.

[5] M. Dorigo, M. Birattari and T. Stutzle, "Ant colony optimization," *Computational Intelligence Magazine, IEEE*, vol. 1, no. 4, pp. 28-39, 2006.

[6] M. Nagy, Z. Akos, D. Biro and T. Vicsek, "Hierarchical group dynamics in pigeon flocks," *Nature*, vol. 464, pp. 890-893, 2010.

[7] V. A. Tucker, A. E. Tucker, K. Akers and J. H. Enderson, Curved ight paths and sideways vision in peregrine falcons (*Falco peregrinus*), *J. Exp. Biol.*, vol. 203, pp. 3755-3763, 2000.

[8] K. S. Galloway. Ph. D. Thesis, Department of Electrical and Computer Engineering, University of Maryland, to appear September 2011.

[9] K. S. Galloway, E.W. Justh and P.S. Krishnaprasad, "Geometry of cyclic pursuit," *Proc. 48th IEEE Conf. Decision and Control*,7485-7490, 2009.

[10] K. S. Galloway, E.W. Justh and P.S. Krishnaprasad, "Cyclic pursuit in three dimensions," *Proc. 49th IEEE Conf. Decision and Control*, 7141-7146, 2010.

[11] K. S. Galloway, E.W. Justh, and P.S. Krishnaprasad, "Portraits of cyclic pursuit," accepted for publication, *Proc. 50th IEEE Conf. on Decision and Control*, 2011.

[12] R. L. Bishop, "There is more than one way to frame a curve," *The American Mathematical Monthly*, vol. 82, no. 3, pp. 246-251, 1975.