What Types of Interactions do Bio-Inspired Robot Swarms and Flocks Afford a Human?

Michael A. Goodrich, Brian Pendleton, and Sean Kerman
Computer Science Department
Brigham Young University
Provo, Utah USA
email: mike@cs.byu.edu

P.B. Sujit
Electrical and Computer Engineering Department
University of Porto
Portugal
email: suit@fe.up.pt

***Abstract*—This paper uses simulations to identify what types of human influence are afforded by the flocking and swarming structures that emerge from Couzin's bio-inspired model [4]. The goal is to allow a human to influence a decentralized agent collective *without resorting to centralized human control*. Evidence is provided that, when nominal agents use switching-based control to respond to human-guided predators and leaders, the resulting behavior is responsive to human input but is obtained at the cost of causing the dynamic structure of the collective to follow a single flocking structure. Leaders are more effective in influencing coherent flocks, but predators can be used to divide the flock into sub-flocks, yielding higher performance on some problems. Introducing a so-called "stakeholder" leadership style makes it possible for a human to guide the agents while maintaining several different types of structures; doing so requires more than one human-controlled agent. We then demonstrate that it is possible to produce potentially useful emergent dynamics without centralized human control, and identify an important type of emergent dynamics: automatic switches between structure types.**

I. Introduction

Steinberg identified human-robot interaction (HRI) with bio-inspired robot teams (BIRT) as an important research area for producing responsive, robust systems for complex surveillance and reconnaissance problems [18]. This paper emphasizes that subset of HRI which helps humans manage multiple remote robots [10, 6]. Similarly, this paper emphasizes that subset of BIRT work which encodes principles of biological societies [20] in robots [16] to produce bio-inspired collective phenomena in the robot teams.

Appropriately combining elements of HRI with BIRT should yield robot teams that can be efficiently managed by humans but that retain robust qualities in the presence of unreliability. Such teams should help break through fan-out barriers imposed by a conventional methods of supervisory control, and help open up the possibilities of large-scale robotic teams being used in surveillance, reconnaissance, fire suppression, and emergency response.

Many bio-inspired teams exhibit robust spatio-temporal structures as attractors induced by decentralized control [21]. When adding human influence to bio-inspired teams, it is imperative that decentralized implementations and distributed communication be preserved. Most existing approaches to human interaction with BIRT use centralized modes of influence such a broadcasting parameter changes to all agents or implementing design-time consensus algorithms to support very specific collective behavior. The work presented in this paper is novel because it allows for decentralized human influence, which may yield systems that are more robust to natural variations or communication dropouts than centralized approaches. Thus, we contribute to a decentralized implementation that allows (a) the human to interact with only a subset of agents, (b) all robots to be responsive to human influence through this subset, and (c) the collective to be robust. In other words, we want the human to be able to influence which collective behavior emerges and then shape how that collective behavior evolves.

This paper does not present a novel robot or user interface design. Rather, the paper evaluates what types of human influences are afforded by bio-inspired robot swarms and flocks. This evaluation uses case studies with low numbers of humans augmented by so-called "Oz of Wizard" studies that simulate the worst-case performance of a human [19]. Since we are measuring what properties of bio-inspired teams afford human interaction rather than human performance per se, it is sufficient to use small numbers of real humans plus some simple simulated human inputs to help identify the impact of various human influence on the behavior of the collective.

The example tasks used in this paper are abstract and simple, but they are sufficient to provide insight into how humans can influence bio-inspired teams. More specifically, the tasks represent a large set of tasks that require coordinated robot movement to a desired location while maintaining a robust collective structure.

II. Related Literature

Miller et al. advocate the use of playbook-style interactions between a human and a team of multiple robots [13], wherein a human calls plays that trigger predictable patterns of behavior. Simple plays, like grouping and searching, have been used to manage several large, simulated teams (50-200 robots) [8]. More complicated plays, like coordinated rendezvous or formation-following, have been applied to smaller teams [11], with some work potentially scaling to large numbers of agents using torus-like behaviors encoded as cyclic pursuit attractors [17]. One simple play is for a

human to control a leader agent who then influences other agents through conventional or bio-inspired means [7].

Other recent work explores centralized methods for influencing patterns of collective behavior so that a human can guide the team [9, 7, 12]. More decentralized approaches include methods to modify potential fields in response to directions from a human operator [2] or virtual agent [15].

Shifting from conventional robot teams to swarm-like teams, Bashyal and Venayagamoorthy [3] presented a "human-swarm interaction" (HSI) approach that provided a human with a partial plan and global information, and then allowed the human to adjust the autonomy of a small subset of swarm members to influence swarm behavior. The GUARDIANS project seeks to use swarm robotic technology to support firefighters [1, 14].

HSI builds on the many formal analyses of bio-inspired and interacting particle systems. Of particular relevance is work that analyzes the emergence of collective structures, including "translating states" and "rotating states" that are very similar to the flocks and toroids used in Couzin's and other work [21]. Others have analyzed how various forms of external influence can shape the way a collective moves through space [15].

III. How Can a Human Influence?

Consider the following general discrete-time model for allowing humans to influence BIRT agents:

$$x^i_{t+1} = f^i(x^i_t, \mathbf{x}^{\neg i}_t) + g^i(x^i_t, u^{\mathrm{op}}_t). \tag{1}$$

In the model, x^i_t represents agent i's state at time t, $\mathbf{x}^{\neg i}_t$ represents the states of every other agent in the collective except for agent i, f^i is the switching controller (defined below) that encodes how the states of other agents affect agent i, and g^i encodes how the operator input, u^{op}_t, influences agent i. A *switching controller* is a nonlinear controller in which the type of response can change depending on the relative values of x^i_t and $\mathbf{x}^{\neg i}_t$; for example, if a predator is close to an agent, then the agent ignores all other agents and flees, but if no predator is close then the agent switches to a controller that causes the agent to try to align to its neighbors.

We assume that g^i depends only on the state of the agent, meaning that agent i does not consider the state of other agents when it computes the influence from the human. Furthermore, we assume that local communication between agents is perfect. Future work should address these simplifying assumptions.

In bio-inspired robot teams, two leadership models have been studied by others: lead-by-attraction and lead-by-repulsion [15]. Lead-by-repulsion is more commonly referred to as *predation* where the leader is a predator and agents are prey, so the leader influences the behavior of the agents by pursuing them. By contrast, lead-by-attraction is often associated with the colloquial use of the word leadership, meaning that a leader is one that gets ahead of a group and the group follows. For simplicity, we will call lead-by-repulsion models *predator-style* and lead-by-attraction models *leader-style*.

Under the predator and leader styles, the human influences leader agents via $g^i_t(x^i_t, u^{\mathrm{op}}_t)$, and then these agents influence other agents via f. In these models, agents are influenced either by the human or by other agents but not by both. Thus, agents will either be leader/predator agents or they will be nominal agents. For leader/predator agents, $f = 0$ and $g \neq 0$; for nominal agents, $f \neq 0$ and $g = 0$.

In 2005, Couzin performed an empirical study that explored how informed individuals can guide the behavior of "animal groups on the move" [5]. This analysis evaluated only the influence of informed individuals on what we identify below as parallel groups. We refer to this type of informed individual as a *stakeholder* because the individual is influenced both by knowledge about the location of a goal and by the behavior of the group; the individual has a stake in reaching the goal and in staying with the collective. Unlike Couzin, we evaluate how stakeholders affect the behavior of more than parallel groups.

A stakeholder tries to sustain influence by remaining in close proximity to the collective, allowing itself to be influenced by the collective as well as to influence it. As before, for nominal agents $f^i \neq 0$ and $g^i = 0$. By contrast to the leader- and predator-style, the stakeholder is influenced both by the human and by the other agents (both $f^i \neq 0$ and $g^i \neq 0$). We discuss the precise implementation of the stakeholder model when we present results for this leadership style.

In the next section, we will show how Couzin's specific model fits the general model from this section, but discussing the general model gives some confidence that the results of this paper may generalize to other systems. Moreover, the distinction between switching and non-switching controllers that are evident from the general model are probably useful for other implementations that use some or all of the zones described in the next section.

IV. Couzin's Model

Couzin's model is a well-studied bio-inspired model that defines three circular zones around an agent: repulsion, orientation, and attraction, with radii of R^{rep}, R^{ori}, and R^{att}, respectively [4]. If there are any other agents in the zone of repulsion, than those agents induce repulsion forces and any agents outside of the zone of repulsion are ignored. If there are no other agents in the zone of repulsion, then any agent within the zone of attraction and/or the zone of orientation induces attraction and/or alignment forces. Couzin's model is thus a nonlinear controller where agents make zone-dependent switches between control laws.

In addition to the radii of the three zones, Couzin's model has four other critical parameters: noise, blindspot angle (ϕ), speed (s), and maximum turning rate (ω). In the interest of space, we only summarize the model as follows: an agent computes a desired direction by being repelled from all agents within the zone of repulsion (or aligned to and attracted to all agents within the zones or orientation/attraction) and then moves toward that desired direction by turning no more than ω radians per second.

We analyze Couzin's model for three reasons. First, it exhibits four qualitatively different behaviors, two of which

have been shown to be attractors in other decentralized systems [21]. The presence of attraction, repulsion, and orientation zones is either typical or is a superset of influences often found in other decentralized systems. Second, the model is compatible with many modern robot systems. The model assumes a constant speed with the only control coming in the form of changes in direction. This closely mimics the qualitative dynamics of unmanned aerial vehicles and is also compatible with four-wheeled robots that must maintain a constant speed. Moreover, sensing and communication between robots is restricted to a constrained communication distance, avoiding power, bandwidth and sensing limitations. Third, Couzin's model is fully decentralized and behaves robustly in the presence of noise. This robustness to noise is essential because the actual dynamics of real robots will naturally have individual variations and will also naturally have different dynamics from Couzin's system.

Simulations in this paper use a sampling rate of $\Delta t = 0.2$ Hz; a noise value sampled uniformly from $[-0.2, 0.2]$ radians is added to angular velocity each second. Following Couzin's example, all spatial units are defined with respect to R^{rep} which is set to unity. Speed is defined with respect to this unit as well, that is, $s = 1$ means that an agent moves one repulsion radius per second. In Couzin's model, forward speed is constant for all agents.

The blindspot needs further explanation. An agent traveling in direction θ can only sense agents that are within $\pm\phi$ radians from the direction of travel. The blindspot is any area outside of $\pm\phi$ radians from the direction of travel. The blindspot opens the possibility of using visual sensing for real robots, but clearly does not go far enough toward modeling real robots by including inevitable things like occlusion.

A. Relating Couzin's Model to Equation (1)

We now show how Couzin's model relates to Equation (1). Couzin's model is a switching controller with dynamics generically given by $x^i_{t+1} = f^i(x^i_t, \mathbf{x}^{\neg i}_t)$. The states are the position vector, $\mathbf{c}^i_t$ and the direction of travel, θ^i_t, yielding $x^i_t = [\mathbf{c}^i_t, \theta^i_i]^T$. The position vector evolves over time as

$$\mathbf{c}^i_{t+1} = \mathbf{c}^i_t + \Delta t \begin{bmatrix} s\cos\theta^i_t \\ s\sin\theta^i_t \end{bmatrix},$$

where s is the speed of the agent. The direction vector evolves by first identifying and then tracking the desired direction of travel, $\hat{\theta}^i_{t+1}$. The desired direction of travel is obtained using the switching control law subject to the blindspot constraints as follows: an agent is repelled by all agents within the radius of repulsion; if no agents are within the radius of repulsion, then the agent is attracted to all agents within the radius of attraction and tries to align with all agents within the radius of alignment. Tracking is then performed using

$$\theta^i_{t+1} = \begin{cases} \hat{\theta}^i_{t+1} & \text{if } |\theta^i_t - \hat{\theta}^i_{t+1}| > \omega\Delta t \\ \text{sign}(\theta^i_t - \hat{\theta}^i_{t+1})(\theta^i_t - \omega\Delta t) & \text{otherwise} \end{cases}.$$

Naturally, this equation assumes appropriate modulo arithmetic in angle space. Importantly, the time delayed coupling of position and angle means that the dynamics for the angle at time $t+1$ depend both on the angle at time t and the angle at time $t-1$. This makes it feasible to believe that the reason that Couzin's model exhibits both a parallel group and a torus phase is because these phases are attractors for a double integrator dynamics model with attractors characterized by the necessary (but not sufficient) conditions $(\dot{\theta} = 0, \ddot{\theta} = 0)$ and $(\dot{\theta} = \alpha, \ddot{\theta} = 0)$ for some constant α, respectively.

B. The Collective Structures from Couzin's Model

Figure 1 illustrates the four phases identified by Couzin; the figures were generated using the parameters shown in Table I. The figures illustrate agents as circles with a line emanating from them in their direction of travel. Each subfigure in Figure 1 is a snapshot at one time instant. The *swarm*

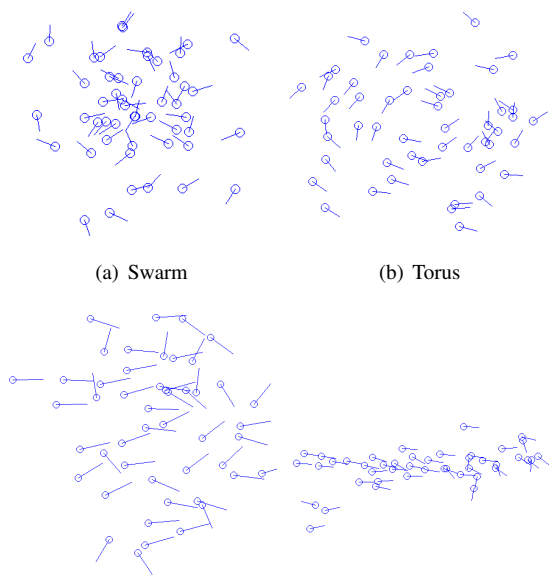

(a) Swarm (b) Torus

(c) Dynamic Parallel (d) Highly Parallel

Fig. 1. Snapshots of various phases from Couzin's model: (a) Swarm. (b) Torus. (c) Dynamic parallel group. (d) Highly parallel group.

phase is a highly dynamic collective structure that tends to stay stationary and exhibits short bursts of agent alignment interspersed with apparently random milling about the swarms center. The *toroid phase* has agents that circle in the same direction around a relatively stationary "hole" in the middle of the structure. The *dynamic and highly parallel* structures both exhibit collective movement in a unified direction, with the highly parallel structure exhibiting greater alignment and greater group velocity.

The title of this paper uses the terms "swarms" and "flocks" as colloquial terms that represent these four types of collective structures. These colloquial terms indicate that collective structure exhibits a high degree of spatio-temporal correlation between agents, in contrast to hive and colony-like behaviors where agents may work collectively but over an apparently decoherent spatio-temporal horizon.

Phase	R^{ori}	ΔR^{att}	s	ω (°/s)	ϕ (°)
Swarm	1	10	3	100	150
Torus	3	12	5	100	150
Dynamic Par.	7	10	3	100	150
Highly Par.	11	15	5	70	120

TABLE I
PARAMETER VALUES USED TO IDENTIFY COUZIN'S PHASES FOR 50 AGENTS. $R^{\text{att}} = R^{\text{ori}} + \Delta R^{\text{att}}$.

Couzin used group polarization and group angular momentum as a useful way to understand and detect the collective phases. Toroids spin and stay in place, so they have high angular momentum and low polarization; individual parallel groups tend to align and travel in a straight line together, so they have low angular momentum and high polarization.

V. SWITCHING CONTROLLERS

Couzin's model requires agents to switch controllers depending on whether there are other agents within R^{rep}, R^{ori}, or R^{att}. In this section, we add another controller to which the nominal agents switch depending on whether they are near a predator or leader. We included the phrase "switching controllers" in the title of this section to emphasize that agents change their behavior whenever they are near a leader (within R^{lead}) or predator (within R^{pred}). In the next two sections, we explore an alternative approach.

The experiments presented in this section only consider how parallel groups with parameters $(R^{\text{rep}}, R^{\text{ori}}, R^{\text{att}}, \phi) = (1, 14, 14, 45°)$ can be augmented with the new switching controller to respond to leaders and predators. This restriction is natural because the switching behaviors in the presence of a leader or a predator cause swarms and toroids to act as if they were parallel phases. Simply put, adding a new switching controller and human influence causes all phases to act like parallel phases, at least while human input is present.

The human influences the collective using either a leader or a predator. As described above, adding human influence to Couzin's model sets $g^i = 0$ for nominal agents and $f^i = 0$ for leaders/predators. The state of nominal agents is augmented to include the type of agent (nominal or leader/predator), and the switching control law for nominal agents is augmented as described in the following sections.

Using a predator or leader allows a human to have centralized[1] locus of human influence, namely the predator or leader, while allowing collective behavior to emerge in a decentralized fashion as a leader influences a handful of nominal agents, those nominal agents affect a handful of others, and so on.

A. Predator Management

Predator-style influence is created from the original model by causing agents (illustrated as fish) to be repelled by a predator if within $R^{\text{pred}} = 30$. The predator moves slightly faster than the fish and can turn much more sharply.

[1]We assume that the human can see all agents. This limits results because, in practice, a human may only know the locations of a small subset of agents.

If a predator nears a nominal agent then that nominal agent is repelled by the predator. When the radius of attraction is greater than the predator's radius, the fish tend to stay close together even when the predator starts to "chase" them. This is illustrated in Figure 2 which shows a predator (shark) steering a group of fish.

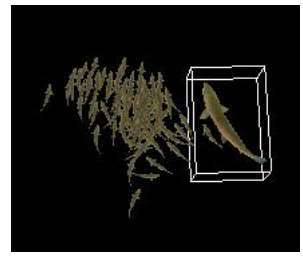

Fig. 2. Using a predator to steer a dynamic parallel group.

B. Leader Management

The leader model is similar to the predator-based model above, but the fish are now attracted to the leader producing a tendency for fish to follow the leader. We set $R^{\text{lead}} = 30$, meaning that the predator-based and leader-based experiments use identical parameters, with the only difference being that predators repel and leaders attract nominal agents.

C. Case Study: Leaders vs. Predators

Differences between predator and leader styles were studied in a scenario with 100 agents (fish) placed in a 120×120 area. Quantities of food, represented graphically as barrels in the simulation, are placed around the map to represent the information to be gathered. The "food" is depleted at 1 unit per second per fish whenever a fish is within $\sqrt{10}$ units, that is, the initial quantity of food is measured in fish-seconds. For example, if N fish are within range of food for t seconds, then $N \times t$ food units are depleted. In the simulations, $R^{\text{rep}} < R^{\text{pred}}$ and $R^{\text{rep}} < R^{\text{lead}}$ meaning that agents follow the leader or are repelled by the predator unless the agents are about to collide with each other. The fish were not aware of the food unless they were in proximity to the food, but the human could see all food sources; thus, fish could consume food if nearby, but could not autonomously find food without human influence.

We performed case studies using both a leader and a predator that were (a) controlled by a single human and (b) controlled by a simulated human that used a worst-case zig-zag path in an "Oz of Wizard" style experiment. Using a convenience sample of six student volunteers, we obtained informed consent and instructed the participants to guide the fish in such a way that the food was consumed as quickly as possible. In the first phase of the experiment, participants were more effective using the predator than the leader because they used the predator to divide the 100 fish into different groups and distribute them around the environment. Leader-style influence didn't allow this divide and conquer approach so performance was lower. In the second phase of the experiment, participants were instructed to maintain a cohesive group.

For each nominal agent, we can determine which other nominal agents are within one of the radii of attraction[2]. We can then define $\mathcal{A}$ as the histogram of the resulting interagent topology over 120 time steps.

The zig-zag path simulates a type of worst-case human behavior wherein the leader/predator ignores the fish and instead shifts direction left and right while maintaining a constant average bearing. This is designed to minimize how long a human sustains influence over any particular subset of agents because the faster leader/predator shifts from one neighborhood of fish to another with each zig. There is a significant discrepancy in the $\mathcal{A}$ histograms when agents use autonomous *zig-zag control* in phase 2. Recall that the zig-zag controller is a type of worst case operator input, because the leader/predator moves in a zig-zag pattern that disrupts the nominal topology as much as possible. Results, shown in

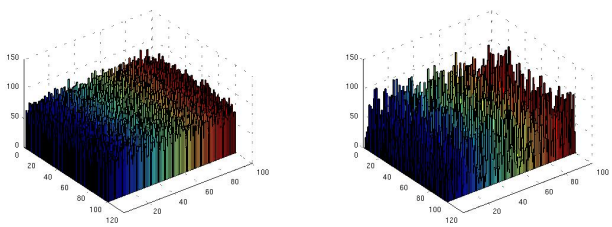

(a) Leader (b) Predator

Fig. 3. $\mathcal{A}$ using the zig-zag controller for (a) leader model and (b) predator model. The values on the axes are not particularly important since the amplitude is a function of the simulation duration, but it is important to note that the same scale is used for both plots.

Figure 3, suggest that leader-based influence maintains cohesion better than predator-based influence when the operator behaves in a somewhat worst-case manner.

Under *human influence* in phase 2, $\mathcal{A}$ was approximately uniform and approximately 120 units tall for both leader-based and predator-based control. This suggests that humans are clever enough to guide agents into a cohesive parallel group phase; the histograms for both leader-based and predator-based influence are very similar to the one shown in Figure 3(a).

To better understand this effect, we measured how long the leader and predator sustained influence over other agents for the case when a human controlled the leader/predator in phase 2 of the experiment. Results for 15 agents are $(m = 47.9, s = 70.6)$ and $(m = 42.2, s = 36)$ for leader and predator, respectively, where the sample mean values are units of food consumed during a fixed-time experiment. Results for 100 agents are $(m = 53.8, s = 53.1)$ and $(m = 53.7, s = 67)$, respectively. No participant was able to keep the 100 agents together using the predator, which accounts for the higher performance (inadvertent divide and conquer). Participants were able to keep 15 agents together with resulting higher performance for leaders. All participants were more confident using leaders than using predators.

For each agent, we created a time series of the number of new agents the leader interacted with plus the number of agents the leader no longer interacted with at each time index. We then computed the power spectral densities for these time series and averaged these together across the six participants; see Figure 4 which shows only results for 100 agents (results for 15 agents are similar). Simply put, predators cause agents to scramble (as evidenced by more frequent switching), making it more difficult to sustain influence even when the human is trying to keep the group together. Results from the zig-zag do not add any insight so they are omitted.

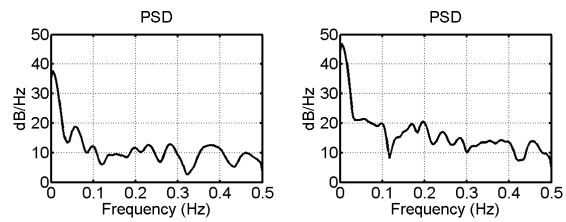

(a) Leader (b) Predator

Fig. 4. Typical power spectral densities for human influence styles.

D. Summary

Results from this section indicate that for parallel group phases, collective structure can be maintained by either leader- or predator-style interaction but the human must be clever to maintain the structure. Simply put, switching controllers afford leader-based human influence when all agents must be guided to the same point, but when it is useful to divide agents then switching controllers afford predator-based influence.

VI. Non-Switching Controllers

Switching controllers tend to preclude anything but parallel group structures. In this section, we evaluate what forms of human influence are afforded by non-switching controllers. We consider two types of non-switching controllers: one based on leadership and one that implements a stakeholder.

Under the leadership style, leader agents respond to human input and ignore nominal agents ($f^i \neq 0, g^i = 0$); nominal agents ignore the human and add the leaders' influences to the influence from every other agent ($f^i = 0, g^i \neq 0$). Under the stakeholder style, nominal agents behave in the same way ignoring the human and adding stakeholders' influences to the influence of every other agent ($f^i = 0, g^i \neq 0$); stakeholder agents are *influenced both by the human and the the nominal agents* ($f_i \neq 0, g^i \neq 0$). This is very similar to Couzin's use of informed individuals that were drawn to a goal but were also governed by other inter-agent dynamics [5].

We implement the stakeholder model by determining both the desired direction using Couzin's equations as well as the direction to a fixed food source placed by a human (a type of waypoint). Couzin's model produces a desired direction $\hat{\theta}_i(t + \Delta t)$ computed from $f^i(x_t^i, \mathbf{x}_t^i)$. This is then added to

[2]We constructed the connectivity matrix assuming that an agent was connected to any other within attraction or repulsion range regardless of whether either agent was actually ignoring the other because of the presence of the leader/predator.

$\hat{\theta}_i^{\text{food}}(t+\Delta t)$, which is the influence caused by the human telling the stakeholder to be attracted to a point of interest through $g^i(x_t^i, u_t^{\text{op}})$; in this case, u_t^{op} is the location of the food source. Adding f to g yields $\tilde{\theta}_i(t+\Delta t) = \hat{\theta}_i(t+\Delta t) + 0.8\hat{\theta}_i^{\text{food}}(t+\Delta t)$ where the 0.8 weighting is subjectively chosen to encourage the stakeholder to stay near the group more than going to the goal.

We conducted an experiment designed to evaluate if and when a stakeholder could influence the different collective phases of Couzin's model, and to compare a stakeholder's influence to that of a leader. To do this, we used the parameters given in Table I. We then use zero or more "Oz of Wizard"-style stakeholders and zero or more leaders of various speeds to see how the different combinations influenced the position of the group.

A. No Leaders

As a point of reference, it is useful to show the final positions of the centroids of each collective phase after 45 seconds in the absence of any type of leader. This is shown in Figure 5. The legend indicates which symbols correspond to

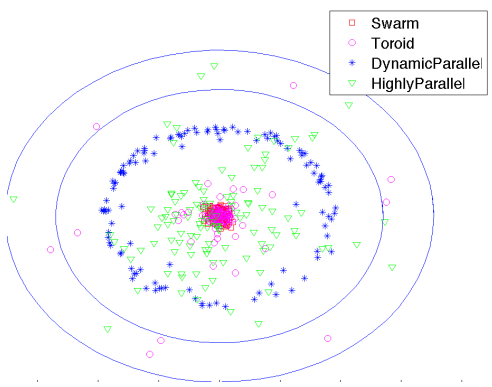

Fig. 5. Distribution of centroids for various phases with no leaders. The mean is approximately $(0,0)$ and $\sigma = (43.05, 43.31)$.

which phase. Ideal agent speed for the various phases differs, so results are shown for agents traveling at speed $s = 3$ for swarms and dynamic parallel groups and $s = 5$ for toroids and highly parallel groups. The inner circle represents the maximum distance that an agent could travel at speed $s = 3$ over the $45s$ simulation duration, and the outer circle is the maximum distance that a torus could travel when agents travel at speed $s = 5$. The radius of the outer circle is $\frac{\pi}{4}s$, which can be shown to be the maximum distance that a torus can travel when agents travel at speed s.

It is useful to give a short description of the collective movement patterns of each phase in the absence of any type of leader of stakeholder.

- Swarms tend to stay closely coupled together and rarely move far from the origin.
- Toroids also tend to stay closely coupled together and rarely move far from the origin. Occasionally, a torus will take some time before it forms, causing it to locate far from the origin.
- Dynamic parallel groups quickly cohere and then start traveling in a somewhat direct line in some random direction. The blue triangles clustered in a circle around the origin suggest that these groups all travel approximately the same straight line distance from the origin.
- Highly parallel groups quickly cohere, but rarely travel in a straight line from the origin, causing a distribution that is within the "maximum distance" circle.

B. Leaders and Stakeholders

Fifty agents were placed in the world and given parameters that caused them to form one of the four phases. Agents were initially distributed randomly around the origin with random initial directions. One or more agents in the group were selected to be leaders or stakeholders and then given information about a stationary food source placed at location $(30, 30)$ in the world. Leader agents go directly toward the food, but stakeholder agents try to reach the food while maintaining their role in the group. Qualitatively, a stakeholder tends to occupy a location in a collective structure that is nearest the food. For example, the stakeholder joins the torus (when the group is in the torus phase) but tends to spend most of its time on the side of the torus nearest to the food. Similar behavior is observed for the swarm and parallel phases.

For the purposes of discussion, we will call leaders and stakeholders by the generic term "manager." The speed of the manager is determined by taking the speed of the nominal agents from Table I and multiplying by the speed scaling factors from the second line of Table II. For example, the speed scaling factor of 0 means that the leaders/stakeholders are stationary, and the speed scaling factor of 1 means that the leaders/stakeholders are traveling at the same speed as the other agents. The first row in Table II indicates the number of agents who are informed of the food source, ranging from 2% to 50% of the total agents.

Number of Leaders/StakeHolders	$\{1, 2, 3, 5, 10, 15, 25\}$
Speed scaling factors	$\{\frac{1}{5}, \frac{1}{3}, \frac{1}{2}, 1, 2\}$

TABLE II
PARAMETERS USED TO CREATE THE SCATTERPLOTS.

The centroids for the different groups are computed as a sum of the locations of each agent, including the locations of the leaders and stakeholders. Since the stakeholders are part of the collective, they contribute to the ultimate location of the centroid and should be used in computing the centroid of the collective at the end of the 40 second simulation. The leaders are not technically part of the collective since they ignore other agents, but we include their final location in the computation of the centroid so as to make sure that there is no unfair advantage given to the stakeholders in comparing scatterplots of stakeholders and leaders.

For both manager styles, one and two managers had only a modest impact on the ultimate distribution of the phase, but as the percentage of managers grows their influence also grows. Figure 6 shows how the error decreases much more rapidly for

stakeholders than leaders, and also that the apparent plateau of error is smaller for stakeholders than leaders.

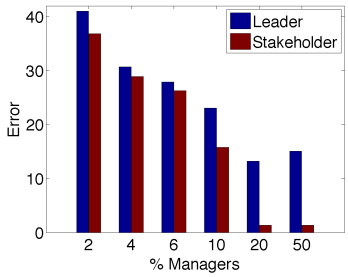

Fig. 6. Error as a function of the percentage of number of managers. Error is Euclidean distance between the final centroid position and the food source.

Once the number of leaders is five or more, constituting at least 10% of the number of agents in the collective, we start to see a difference in influence for stakeholders over leaders, even given the fact that centroids for leaders should be biased more toward the food location since all leaders end up at that location.

A sample of results is given in Figure 7 which illustrates the final distribution of collective centroids when there are ten managers. The scatterplots are results for all speed scaling factors. Compared to leaders, stakeholders yield (a) centroid locations that are closer to the food source and (b) smaller standard deviations of centroids over all the trials. To emphasize the differences in the means, two solid lines intersect at the location of the food sources and two dashed lines intersect at the location of the mean of the centroids. These results extend Couzin's results [5] by showing that (a) large numbers of informed individuals have more influence over a group than small numbers, (b) informed individuals can influence more than just dynamic parallel groups and (c) torus and flock attractors can be guided, causing toruses to translate and influencing flocks to travel in a desired direction.

We conclude that non-switching controllers afford stakeholder-style influence, at least when it is important for the agents to maintain their collective structures. Moreover, regardless of the management style, multiple manager agents are required to produce satisfactory influence.

C. Emergent Influence Dynamics: Structure Switching

In the previous section, the radius of attraction was low enough that the collective structures sometimes fragmented, leaving groups of agents that were not within each others radii of influence. In this section, we explore what happens when we guarantee that agents stay coherent by making the radius of attraction infinite and by eliminating the blind spot. This allows us to see what happens when the influence of leaders is sufficiently high to affect the behavior of all agents to some extent. Interestingly, potentially useful group behaviors emerge under these conditions.

Let $s_\ell = 1$ denote the speed of the leader and consider what happens when the speed of the nominal agents is faster than the leader. Set the parameters of Couzin's model to $R^{\text{rep}} = 1$,

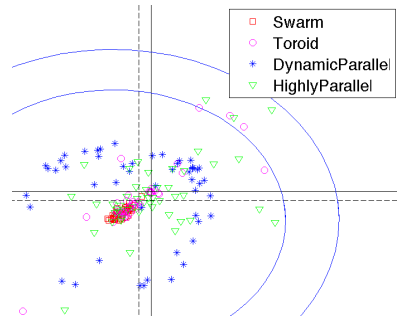

(a) Leader

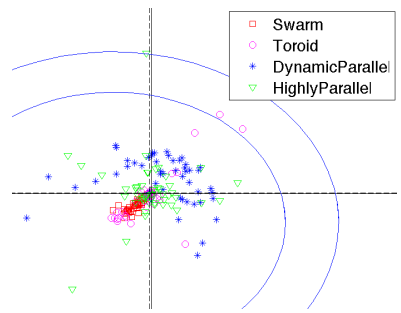

(b) Stakeholder

Fig. 7. 20% of the agents try to influence the collective: (a) ten leaders ($m = 20.56, 20.79$), $\sigma = (33.03, 33.21)$); (b) ten stakeholders ($m = (28.77, 30.59)$, $\sigma = (23.16, 24.37)$). The solid lines denote the location of the food, and the dashed lines the location of the mean of all types.

and $R^{\text{ori}} = 3$ so that a toroid structure emerges. Consider a group of $N = 80$ agents, and consider groups of 30 and 50 leaders. In this set of simulations, leaders are added to the nominal group so, for example, when a group of $N = 80$ has 30 leaders there are a total of 110 agents. The "Oz of Wizard" leaders travel due east at constant speed, but travel more slowly than the centroid of the torus. Note that leaders are excluded from any calculations of centroid and moment.

When agent speeds are moderately higher than the speed of the leaders ($s = 10$ and $s = 5$), the torus tends to follow the leaders but had difficulty keeping up with it (this is true for $\omega \in \{40, 70, 100\}^\circ/sec$). When the torus got behind the leaders, it would break from the torus formation, temporarily form a dynamic parallel group, catch up with the leader, and then resume the torus.

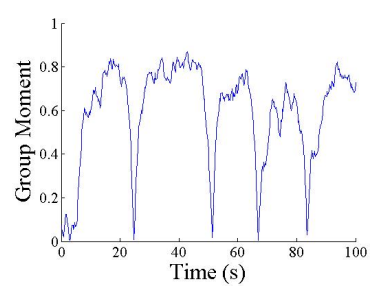

Fig. 8. When the distance of the group gets large, the group changes from a torus to a parallel group.

Figure 8 illustrates this emergent switching behavior, where the switching occurs at the level of the collective structure rather than at the individual. We choose to illustrate this switching using group angular momentum, since angular momentum is very high for a torus and very low for a parallel group. Although the individual agents are not programmed to switch behaviors when the leaders get too far away, the collective structure of the agents automatically changes when the leaders get too far away. The figure illustrates this in a series of changes in group angular momentum that correspond to large distances from the centroid of the toroid and the locations of the leaders.

This natural emergence of switches in collective structure could be very useful for the design of systems that enable a human to influence and manage a large team of robots. Simply put, emergent switches between collective structures could afford a human the flexibility to naturally manage collective phases without resorting to explicit centralized control. Moreover, this switching likely indicates that it should be possible for a human to change which attractor (torus or dynamic parallel group) a particular group exhibits without changing internal agent parameters; since the attractors tend to be robust to external disturbance, switching between attractors means that the human's influence can persist until a sufficient disturbance causes the agents to leave the attractor.

VII. Conclusions

Many bio-inspired models, such as Couzin's, can produce a variety of collective structures, but much prior work on influencing these models require centralized forms of human influence. We have used small case studies and "Oz of Wizard"-style experiments to explore how a human can influence these collective structures using decentralized methods. Augmenting the switching controllers with a new switch that responds to the presence of a human-controlled predator or leader allowed a human to guide the collective, but only for one type of collective structure. Moreover, the leader style afforded more sustainable influence over the collective and was compatible with leading the collective to a desired location, and the predator-style afforded divide-and-conquer approaches to problem solving.

Non-switching controllers can be used to allow a human to influence a broader set of collective structures, but more than one human-influenced manager is needed. Stakeholder-style management in which the manager both sought to influence and to be influenced afforded more human control over the collective phases than a management style that did not allow the leader to be influenced by the group. Importantly, the use of decentralized leaders can be used to induce emergent phenomena that may afford use by humans. These emergent phenomena include the ability to switch from one structure to another and back without changing the parameters of the agents. Future work should explore how decentralized human influence can enable a team of distributed humans to manage a bio-inspired collective.

Acknowledgments

This work was funded by the Office of Naval Research and the Army Research Lab. Opinions do not necessarily reflect the opinions of the funding agencies.

References

[1] L. Alboul, J. Saez-Pons, and J. Penders. Mixed human-robot team navigation in the GUARDIANS project. In *Proc. of the Intl. Conf. on Safety, Security, and Rescue Robotics*, Sendei, Japan, October 2008.

[2] L. Barnes, M. Fields, and K. Valvanis. Umanned ground vehicle swarm formation control using potential fields. In *Proc, of the Mediterranean Conf. on Control & Automation*, Athens, Greece, July 2007.

[3] S. Bashyal and G. K. Benayagamoorthy. Human swarm interaction for radiation source search and localization. In *IEEE Swarm Intelligence Symposium*, St. Lous, Missouri, USA, 2008.

[4] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and H. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1), September 2002.

[5] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433:513–516, 2005.

[6] M.L. Cummings, S. Bruni, S. Mercier, and P.J. Mitchell. Automation architecture for single operator, multiple UAV command and control. *The International C2 Journal*, 1(2):1–24, 2007.

[7] X. C. Ding, M. Powers, M. Egerstedt, S. Young, and T. Balch. Executive decision support: Single agent control of multiple UAVs. *IEEE Robotics and Automation Magazine*, 2009.

[8] B. Hardin and M. A. Goodrich. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *Proc. of ACM/IEEE Intl. Conf. on Human-Robot Interaction*, March 2009.

[9] Z. Kira and M. A. Potter. Exerting human control over decentralized robot swarms. In *Proc. of the Intl. Conf. on Autonomous Robots and Agents*, pages 566–571, 2009.

[10] M. Lewis, H. Wang, S. Chien, P. Scerri, P. Velagapudi, K. Sycara, and B. Kane. Teams organization and performance in multi-human/multi-robot teams. In *Proc. of the IEEE Intl. Conf. on Systems, Man, and Cybernetics*, Istanbul, October 2010.

[11] T. McLain and R. W. Beard. Cooperative rendezvous of multiple unmanned air vehicles. In *AIAA Guidance, Navigation and Control Conference*, Denver, Colorado, USA, 2000.

[12] J. McLurkin, J. Smith, J. Frankel, D. Sotkowitz, D. Blau, and B. Schmidt. Speaking swarmish: Human-robot interface design for large swarms of autonomous mobile robots. In *Proc, of AAAI Spring Symposium*, Stanford, CA, USA, 2006.

[13] C. A. Miller, H. B. Funk, M. Dorneich, and S. D. Whitlow. A playbook interface for mixed initiative control of multiple unmanned vehicle teams. In *Proc. of the 21st Digital Avionics Systems Conf.*, volume 2, pages 7E4–1 – 7E4–13, November 2002.

[14] A. M. Naghsh, J. Gancet, A. Tanoto, and C. Roast. Analysis and design of human-robot swarm interaction in firefighting. In *Proc. of the Intl. Symp. on Robot and Human-Interactive Communication*, Muhich, Germany, August 2008.

[15] R. Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.

[16] E. Sahin. Swarm robotics: From sources of inspiration to domains of application. In *Lecture Notes in Computer Science*, volume 3342. SpringerLink, 2005.

[17] A. Sinha and D. Ghose. Generalization of linear cyclic pursuit with application to rendezvous of multiple autonomous agents. *IEEE Transactions on Automatic Control*, 51(11):1819–1824, 2006.

[18] M. L. Steinberg. Biologically-inspired approaches for self-organization, adaptation, and collaboration of heterogeneous autonomous systems. In *Proc. of SPIE Volume*, volume 8062, April 2011.

[19] A. Steinfeld, O. C. Jenkins, and B. Scassellati. The "Oz of Wizard": Simulation the human for interaction research. In *Proc. of the ACM/IEEE Intl. Conf. on Human-Robot Interaction*, 2009.

[20] D. J. T. Sumpter. The principles of collective animal behavior. *Philosophical Transactions of the Royal Society B*, 361:5–22, November 2006.

[21] L. Mier y Teran-Romero, E. Forgoston, and I. B. Schwartz. Noise, bifurcations, and modeling of interacting particle systems. Technical Report Rep US 2011, NIH, 2011.