

Human-Swarm Interactions Based on Managing Attractors

Daniel S. Brown
AFRL Information Directorate
Rome, NY 13441
daniel.brown.81@us.af.mil

Sean C. Kerman
Brigham Young University
Provo, UT 84602
sean.kerman@gmail.com

Michael A. Goodrich
Brigham Young University
Provo, UT 84602
mike@cs.byu.edu

ABSTRACT

Leveraging the abilities of multiple affordable robots as a swarm is enticing because of the resulting robustness and emergent behaviors of a swarm. However, because swarms are composed of many different agents, it is difficult for a human to influence the swarm by managing individual agents. Instead, we propose that human influence should focus on (a) managing the higher level attractors of the swarm system and (b) managing trade-offs that appear in mission-relevant performance. We claim that managing attractors theoretically allows a human to abstract the details of individual agents and focus on managing the collective as a whole. Using a swarm model with two attractors, we demonstrate this concept by showing how limited human influence can cause the swarm to switch between attractors. We further claim that using quorum sensing allows a human to manage trade-offs between the scalability of interactions and mitigating the vulnerability of the swarm to agent failures.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*multiagent systems, coherence and coordination*;
H.1.2 [Models and Principles]: User/Machine Systems—*human factors*;
I.2.9 [Artificial Intelligence]: Robotics—*operator interfaces, intelligent vehicles*

Keywords

Human-swarm interaction, managing attractors of dynamic systems, bio-inspired swarms, quorum sensing

1. INTRODUCTION

Swarms provide complex behaviors out of simple agents following simple rules. The ability for swarms to adapt and produce complex behaviors allows us to build multi-agent systems that are flexible and resilient. However, two important problems emerge when we try to allow humans to manage robot swarms. First, in contested or noisy real-world

environments where bandwidth is limited, it is necessary to interact with only a subset of the agents in order to achieve desired behaviors. Thus operators need ways to manage a swarm by managing the collective emergent behaviors while only requiring communication with a subset of the swarm. Second, as swarms grow in size it becomes increasingly difficult for a human operator to manage the behavior of a swarm by controlling individual agents.

The swarm model presented in [17] has three properties that make it ideal for studying human interaction with robot swarms. First, the model (a) uses an abstract but appropriate model for how well a robot can sense neighboring robots and (b) matches the mobility characteristics of fixed-wing unmanned aerial vehicles and a large class of fixed-speed non-holonomic robots. Second, the model has been shown to have two emergent collective behaviors: a flock and a torus. Third, the model can be extended, as shown below, to include bio-inspired quorum-signaling, which lends itself to human interaction through managing trade-offs between scalability and vulnerability of the swarm. Note that this property is essential for this paper because the results do not include experiments with real humans or real robots.

The second property is essential for human-swarm interaction because it allows a huge number of details to be ignored when managing the swarm. The behavior of any individual agent in a swarm can be difficult to understand and explain due to the highly dynamic interactions with neighboring agents, but by managing the collective as a whole, much of this difficulty can be avoided. We empirically show that there exists a critical point in the parameter space and show that managing the swarm at this critical point allows both a flock and a torus to form. Through empirical simulations, we show that controlled switches between a torus and a flock are possible, given a significant number of human-influenced agents in the swarm.

The third property is also essential for human interaction with robot swarms because it is necessary to balance the desire to allow the swarm to be responsive to human input without becoming a centralized system with a single point of failure. To increase the scalability of the swarm we introduce the notion of quorum sensing, as found in biological systems and show how this can be applied to a swarm. In addition to increasing the scalability of human-swarm interactions, we show evidence that quorum sensing provides a mechanism for limiting the vulnerability of the swarm to agent failure.

The main contribution of this paper is a theoretical and empirical exploration of how human interaction with robot swarms can operate at a high level of abstraction. Claims

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '14, March 3–6, 2014, Bielefeld, Germany.

Copyright 2014 ACM 978-1-4503-2658-2/14/03 ...\$15.00.

<http://dx.doi.org/10.1145/2559636.2559661>.

in the paper are supported by so-called ‘‘Oz of Wizard’’ empirical studies [28] that explore how an idealized human can manage swarms by managing attractors and trade-offs.

2. SWARM MODEL

The experiments and simulations in this paper build upon the model of swarming presented in [17]. This model exhibits both a flock and torus behavior, is similar to many biological models of swarming behavior, and has dynamics similar to those of actual robots. The model consists of a set of N agents with the dynamics for agent i given by

$$\dot{x}_i = s \cdot \cos \theta_i, \quad \dot{y}_i = s \cdot \sin \theta_i, \quad \dot{\theta}_i = \omega_i \quad (1)$$

where $[x_i, y_i]^T \in \mathbb{R}^2$ is the agent’s position, $\theta_i \in [-\pi, \pi]$ is the agent’s angular heading, s is the constant agent speed, and w_i is the agent’s angular velocity. For simplicity we define $v_i = [\cos(\theta_i), \sin(\theta_i)]^T$ and $c_i = [x, y]^T$. Let $A(t) = a_{ij}(t)$ denote the sensory adjacency matrix where $a_{ij}(t) = 1$ means that agent j is visible to agent i at time t . Each $a_{ij}(t)$ is determined at time t according to a Bernoulli random variable with parameter $p_{ij}(t) = \min(1, 1/d_{ij}(t))$ where $d_{ij}(t)$ is the Euclidean distance between agents i and j at time t .

Similar to the Couzin model of biological swarms [9] and the Reynold model of synthetic agents [26], agents in this model react to neighbors within three different zones: repulsion, orientation, and attraction. The neighbors in these zones are determined by

$$n_i^r = \{j : \|c_i - c_j\|_2 \leq R_r, a_{ij} = 1\} \quad (2)$$

$$n_i^o = \{j : \|c_i - c_j\|_2 \leq R_o, a_{ij} = 1\} \quad (3)$$

$$n_i^a = \{j : a_{ij} = 1\} \quad (4)$$

where n_i^r , n_i^o , and n_i^a are the sets of agent i ’s neighbors in the regions of repulsion, orientation, and attraction, respectively. The parameters R_r and R_o are the associated radii of repulsion and orientation. The angular velocity ω_i is determined by first computing the repulsion, orientation, and attraction vectors

$$u_i^r = - \sum_{n_i^r} \frac{c_j - c_i}{\|c_j - c_i\|_2^2} \quad (5)$$

$$u_i^o = \frac{v_i + \sum_{n_i^o} v_j}{\|v_i + \sum_{n_i^o} v_j\|_2} \quad (6)$$

$$u_i^a = \frac{\sum_{n_i^a} (c_j - c_i)}{\|\sum_{n_i^a} (c_j - c_i)\|_2}. \quad (7)$$

Next, the desired heading vector u_i is computed as $u_i = u_i^r + u_i^o + u_i^a$. Finally, angular velocity, ω_i , is computed as

$$\omega_i = k(\text{atan2}(u_i^y, u_i^x) - \theta_i) \quad (8)$$

where k is a positive gain and $\text{atan2}(u_i^y, u_i^x)$ is the two-argument arctangent that places the angle in the correct quadrant by considering the y and x components of u_i .

We note three important aspects of this model that give confidence that the results will extend to real robots. First, because we limit $(\text{atan2}(u_i^y, u_i^x) - \theta_i)$ to the interval $[-\pi, \pi]$, the magnitude of ω_i is bounded by $k\pi$. Thus the agent dynamics match the Dubins curve model which is often used for actual UAV path planning and applies to some constant-speed, non-holonomic ground robots [10], lending some ecological validity to our experimental results below. Second, the method of choosing neighbors is similar to the random




Figure 1: The two group types formed by our model. Agent headings are represented by straight lines.

neighbor model used in [1] which replicated field observations of starlings and is relevant for robot systems where visibility and sensing distance-limited. Third, we note that the performance in this swarm model is robust to small deviations in model parameters and noise levels [16], indicating that results are likely to be applicable on real robots.

2.1 Group metrics

In order to define the two different attractors of the model we use two metrics of group behavior, namely, group angular momentum, m_{group} , and group polarization, p_{group} [9, 17]. The m_{group} is a measure of the degree of rotation of the group about its centroid and is a value between 0 and 1. The m_{group} of a swarm reaches a maximum value of 1 if all the agents are rotating around the group centroid in the same direction. The p_{group} measures the degree of alignment among individuals within the group and is also a value between 0 and 1. The p_{group} of a swarm reaches a maximum value of 1 when all the agents have the same heading.

2.2 Group types

When $R_o > R_r$, and agents are started from random initial conditions, simulation results only exhibit two group types: a torus and a flock (see Figure 1). A stable torus is characterized by p_{group} close to 0 and m_{group} close to 1. It has a relatively stationary group centroid, and either a clockwise or counterclockwise rotation. A swarm in the torus formation could be used for omnidirectional sensing of a target, perimeter monitoring, or as a fixed-wing UAV loitering command.

A stable flock is characterized by p_{group} close to 1 and m_{group} close to 0. It has a moving centroid with all of the agents heading in the same general direction. A swarm in the flock formation allows a swarm to quickly move from one location to another and could be used for search or tracking.

For a stable flock and torus, m_{group} and p_{group} are inversely related; however, when the swarm is perturbed, these values may fluctuate widely. To enable classification of swarm behaviors despite human influence, we define the group type of the swarm as follows:

$$type = \begin{cases} torus, & \text{if } p_{group} < 0.25 \text{ and } 0.75 < m_{group} \\ flock, & \text{if } m_{group} < 0.25 \text{ and } 0.75 < p_{group}. \end{cases} \quad (9)$$

These values were chosen to ensure the fundamental characteristics mentioned previously were visually evident.¹

¹Unless otherwise noted, $N = 100$, $s = 5$, $k = 0.5$, $R_o = 8$, and $R_r = 1$. Simulations used a discrete time approximation with simulation time step of $\Delta T = 0.1$ seconds.




Figure 2: The group momentum and polarization as the radius of orientation is increased and decreased.

3. ATTRACTORS AND HYSTERESIS

Kerman et al. [17] show that the torus behavior is a fundamental attractor of the swarm model presented above. In this section we briefly summarize work that is under review for a journal that argues that the flock behavior is also an attractor of the swarm model. This attractor is fundamentally caused by the orientation dynamics, with attraction and repulsion simply causing the flock to stay cohesive but avoid collisions. Being able to treat a flock and torus as basic attractors of a swarm system means that a human can treat the torus and flock at an abstract level as cohesive units without having situation awareness of (or even observing [2]) all individual agents.

To argue that the flock is an attractor, consider a discrete time approximation of the system with only orientation dynamics ($u_i = u_i^o$ in Equation (8)). The angular velocity $\dot{\theta}$ can be approximated by

$$\dot{\theta}_i = \frac{k}{n_i + 1} \sum_j (\theta_j - \theta_i) \quad (10)$$

where n_i is the number of neighbors of agent i . When the underlying orientation graph is connected, (10) is known to cause all agents to converge to a common heading [21]. Simulation results show that even when the orientation graph dynamically switches, agents converge to a common heading. Thus, agents following the dynamics in Equation (10) will converge to a stable flock, providing strong evidence that the flock behavior is an attractor of the swarm model.

3.1 Hysteresis

Systems that have inherent memory are said to exhibit hysteresis. Couzin noted hysteresis in his model [9], so there is reason to believe that hysteresis will be found in our model as well. Figure 2 shows the results of incrementally changing the radius of orientation, R_o . To obtain these results, the parameter R_o was incremented or decremented by 1 unit every 1000 time steps, and the average m_{group} and p_{group} was calculated over 15 replicates. As shown in Figure 2, the group behavior depends on the previous history of the group—this is because trajectories in the state space² are characterized by the model parameters, so changing R_o changes the basins of attraction for the flock and torus. Interestingly, when R_o decreases, the flock does not ever switch to a torus but simply remains a flock until $R_o = 0$ at which point the flock turns into an unoriented cyclic group with agents rotating in both directions around the swarm’s centroid.

²The state is the agents’ positions and headings.

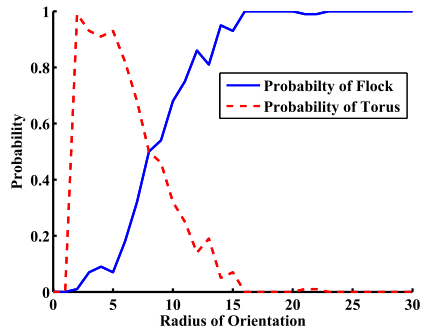


Figure 3: Probability of swarm forming a flock or a torus as a function of radius of orientation.

This form of hysteresis is important because it indicates that the attractors of the swarm exhibit what Couzin called collective memory. This memory is important if the human is going to manage a swarm by managing the attractors for two reasons. First, if the attractors are unstable, the human will not be able to form an accurate and reliable mental model of the swarm’s behavior. Second, if the attractors are stable, the human can neglect the swarm, meaning that the human need not exert continuous control.

Unfortunately, this form of hysteresis is found by increasing or decreasing the parameter R_o for all agents, which requires the human to act as a centralized controller of the agents by broadcasting global parameters. Fortunately, it is apparent from inspecting Figure 2 that there are two different attractors that can be obtained when R_o is roughly in the interval (4, 12). We use this in the next sections.

3.2 Tipping point

Because we desire to be able to switch between the torus and flock attractors without changing model parameters, we ran an experiment to determine whether parameter values exist that allow both group types to emerge. We ran a series of simulations using $N = 100$, $k = 0.5$, $s = 5$, $R_r = 1$ and varied the radius of orientation. Each simulation was run for 200 seconds. One hundred simulations were performed for values of R_o between 0 and 30. For each iteration, agents were given random initial headings with random initial positions uniformly distributed over a 20 unit \times 20 unit square.

The percentage of trials that converged to a torus and to a flock were calculated for each value of R_o and are shown in Figure 3. The value $R_o = 8$ resulted in an approximately equal proportion of torus and flock group types. Thus, $R_o \approx 8$ is the tipping point between the basins of attraction for the torus and flock, meaning that when a swarm is started from random initial conditions, both attractors are approximately equally probable. This equal probability suggests that having the human inject randomization and constraints into the swarm is a sufficient condition for allowing the human to determine what behavior will be exhibited.

The hysteresis exhibited when $R_o = 8$ and the tipping point property at the same parameter value indicate that it is possible for a human to change the collective behavior of the swarm by simply perturbing a subset of the agents enough. For this reason we use $R_o = 8$ for the remainder of this paper, but note that our results are robust to small deviations in this and other model parameters.

4. SWITCHING BETWEEN ATTRACTORS UNDER HUMAN INFLUENCE

In this section we explore two leadership strategies that enable a human to switch between the flock and torus attractors. Because the torus and flock are both stable collective behaviors (within an appropriate region of the parameter space) it is interesting to see how much effort is needed to push the swarm out of one stable attractor and into another. Additionally, having a swarm that can switch between different stable formations affords the operator increased flexibility in using the swarm. For example, switching from a torus to a flock, and then back to a torus, would allow a loitering swarm to transition to a stable flock, quickly move to a new location, and then switch to a stable torus to start monitoring that location. To investigate whether these controlled collective behavior switches can be performed when communication is limited, we assume that the human can only interact with a subset of the swarm. We use a refined version of the *stakeholders* that we have used in prior work [25, 13], modified to support switching between attractors.

4.1 Stakeholders

Stakeholders are agents in the swarm that are influenced by both the human and by other agents. Each stakeholder has a priority parameter $\rho \in [0, 1]$ that determines the priority of human influence over influence from other agents. If ρ is high, then the stakeholder responds more to human commands. If ρ is low, then the stakeholder responds more to its neighbors. Human influence is applied via teleoperation through a waypoint q .

Stakeholders can be led using two different methods: attraction and orientation. Stakeholders that are led by attraction have the desired direction $u_i = u_i^{sa} + u_i^o + u_i^r$ where

$$u_i^{sa} = \frac{\rho \hat{q}_i + (1 - \rho)u_i^a}{\|\rho \hat{q}_i + (1 - \rho)u_i^a\|_2} \quad \text{and} \quad \hat{q}_i = \frac{q - c_i}{\|q - c_i\|_2}. \quad (11)$$

The vector $q \in \mathbb{R}^2$ is the human-generated waypoint, $\rho \in [0, 1]$ is priority parameter, and u_i^a , u_i^o , u_i^r are the usual attraction, orientation, and repulsion influences described previously in Equations (5)–(7). This causes stakeholders to continue to orient with their neighbors while adjusting their positions to be closer to the human-specified waypoint.

Stakeholders that are led by orientation have the desired direction $u_i = u_i^a + u_i^{so} + u_i^r$ where

$$u_i^{so} = \frac{\rho \hat{q}_i + (1 - \rho)u_i^o}{\|\rho \hat{q}_i + (1 - \rho)u_i^o\|_2}. \quad (12)$$

This causes stakeholders to continue to be attracted to their neighbors' positions while adjusting their headings towards the human-specified waypoint. We use the notation M to denote the number of stakeholders.

4.2 Methods

To determine how best to switch the swarm from one group type to the other, we ran simulations with values of M ranging from 10 to 100 in 10 step increments and values of ρ ranging from 0.1 to 1 in 0.1 step increments. Both orientation and attraction leadership strategies were tested using a constant waypoint command. These methods systematically explore the bounds of what an idealized human, which we denote as an OoWiz after the Oz of Wizard approach [28], can do to switch between attractors. Ten trials

were performed for each M and ρ pair, and for each leadership strategy. To initialize the simulations, agent positions were distributed randomly in a 20 unit \times 20 unit square centered at the origin. Initial orientations varied depending on whether the group began as a torus or flock.

When switching from a flock to a torus, all agents were given random initial positions with initial headings $\theta_i = 0$, $\forall i$. After allowing 25 seconds for the group to stabilize, the constant input $\hat{q} = c_g(25) + [0, 10]^T$, where $c_g(25)$ is the group centroid after 25 seconds, was applied to each stakeholder to encourage the group to turn and form a torus. When switching from a torus to a flock, each agent was given a random initial position c_i with initial heading $\theta_i = \text{atan2}(c_i^x, c_i^y) + \pi/2$ to form a counterclockwise torus. After letting the group stabilize for 25 seconds, an arbitrarily large constant control input, $\hat{q} = c_g(25) + [10, 000, 0]^T$, was applied to the stakeholders to influence them to form a flock. We gave the swarm 200 seconds to switch group types, removed the OoWiz human influence, and gave the swarm 50 seconds to stabilize to evaluate whether the swarm would remain in the desired group type.

4.3 Switching from Flock to Torus

Leadership by attraction was effective for switching from a flock to a torus. Figure 4(a) shows the percentage that switched when under the influence of the OoWiz human. Figure 4(b) shows the percentage that switched after the OoWiz human input was removed. As can be seen, for sufficiently high M and ρ , the group successfully switched. However, there is a noticeable drop in the number of simulations that switched and remained for $\rho = 1$ and high values of M . We investigated this and found that because the OoWiz human was explicitly controlling the attraction dynamics, the agents formed a flock-like structure that circled around the reference input. Because the agents never spread out into a full torus, when the control input was removed the group returned to a flock formation.

When leading stakeholders by orientation, we found a large discrepancy between the number of simulations that switched to a torus and the number of simulations that remained as a torus (see figures 4(c) and 4(d)). We investigated this phenomenon and found that the human influence on the orientation dynamics often caused stakeholders to rotate in different directions around the desired torus centroid. This behavior prevents the torus from fully forming and is undesirable for actual robots because of the risk of head-on collisions. After the human influence was removed, the stakeholders were able to orient properly, causing a torus to form. This is an example of what Walker et al. refer to as *neglect benevolence* [31]. Neglect benevolence says that, in some cases, the swarm must be allowed to self stabilize before receiving a new command from the human.

4.4 Switching from Torus to Flock

When switching from a torus to a flock (see Figure 5), leading stakeholders by orientation worked much better than leading stakeholders by attraction. We examined these results and found that leading stakeholders by attraction was successful in causing the agents to switch from a torus to a flock. However, the attraction input \hat{q} caused the stakeholders to slowly pull away from the rest of the group causing the flock to elongate. Thus, when the human influence was removed, the flock was unstable and usually reformed a torus.

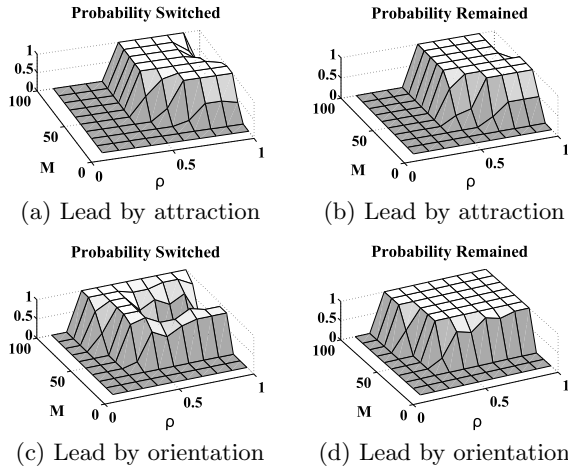


Figure 4: Switching from flock to torus.

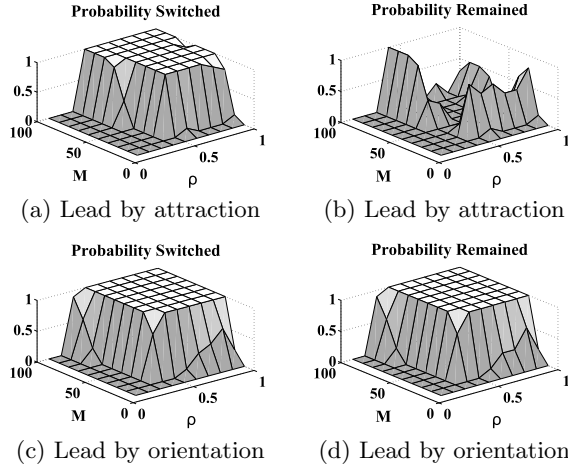


Figure 5: Switching from torus to flock.

Figures 5(a) and 5(b) show that except for limited areas of the parameter space, simulations that switched to a flock usually switched back to a torus. Using leadership by orientation eliminated this phenomenon and caused the agents to form a less elongated flock that remained stable after human influence was removed (see figures 5(c) and 5(d)).

4.5 Discussion

Based on these results, we lead stakeholders by orientation when switching from a torus to a flock, and by attraction when switching from a flock to a torus, for the remainder of this paper. This is a potential problem when viewed in light of our assumption that we want the human to be able to manage a swarm at a high level of abstraction because invoking a switch between attractors requires the human to choose what type of influence to exert. A simple solution to this problem is to allow the human to indicate, through a GUI or some other means, a desire to change from one structure to another. Given the current structure and the desire to switch structures, the correct controller can automatically be used to facilitate the switch.

5. QUORUM SENSING

The previous section showed that an OoWiz human could manage a swarm by switching between attractors and argued that the HRI principles of neglect tolerance and neglect benevolence apply to such a system. Unfortunately, doing so required the human to send information to a large number of robots in the collective. In this section, we propose a way to decrease the number of agents that a human must influence, and find a design parameter that allows a human to manage a trade-off between two important properties of the collective: scalability and vulnerability.

To accomplish this, we use the concept of quorum sensing, which is used in biological systems to regulate the emergence of different behaviors depending on external thresholds [6, 7]. To facilitate a discussion of quorum sensing, we utilize the taxonomy of agent types presented in [12]. An agent’s type determines how it responds to human influence, external influence, and other agents. Agents can be divided into two classes: *human-aware* agents who can respond to human input and *human-blind* agents who do not respond to human input. In this section we explore the properties of a heterogeneous swarm in which we have both human-aware agents and human-blind agents. Specifically, we investigate using stakeholders, a kind of human-aware agent, and a particular kind of human-blind agent called a *type-aware* agent.

5.1 Type-aware agents

Type-aware agents are influenced by both human-aware agents and human-blind agents, but are not influenced by the human. Type-aware agents have an awareness parameter $\alpha \in [0, 1]$ that determines the degree of type-awareness of the agent. If α is high, then the type-aware agent is influenced more by human-aware agents than by other human-blind agents. If α is low, then the type-aware agent is influenced more by other human-blind agents and less by human-aware agents. For the remainder of this section we assume that in the quorum sensing model, agents are either stakeholders or type-aware agents. Thus, for a swarm of size N with M stakeholders, we have $N - M$ type-aware agents.

Type-aware agents follow the normal agent dynamics described in Section 2 unless one or more of their neighbors are stakeholders. Each type-aware agent i uses quorum sensing to adjust its awareness parameter α_i . To incorporate quorum sensing into the agent dynamics, we define a quorum threshold Q_i for type-aware agent i . If type-aware agent i has more than Q_i stakeholder neighbors, then it will temporarily increase its type awareness by setting $\alpha_i = \alpha_i^{max}$, and maintain this increased type awareness until the number of stakeholders within its sensing range falls below Q_i , at which point the agent sets $\alpha_i = \alpha_i^{min}$. We use α_i^{max} and α_i^{min} to denote the maximum and minimum awareness for agent i . The parameter α_i^{max} determines how much agent i is influenced by stakeholder neighbors if it senses a quorum. The parameter α_i^{min} determines how much agent i is influenced by stakeholder neighbors if it does not sense a quorum. Thus, if we let $\mathcal{N}_i = \{j : a_{ij} = 1\}$ represent the set of neighbors of agent i , and \mathcal{S} be the set of stakeholders in the swarm, then for any type-aware agent i we have

$$\alpha_i = \begin{cases} \alpha_i^{max}, & \text{if } Q_i < |\mathcal{N}_i \cap \mathcal{S}| \\ \alpha_i^{min}, & \text{if } |\mathcal{N}_i \cap \mathcal{S}| \leq Q_i. \end{cases} \quad (13)$$

The desired direction vector of type-aware agent i is

$$u_i = \frac{u_i^{ta}}{\|u_i^{ta}\|_2} + \frac{u_i^{to}}{\|u_i^{to}\|_2} + u_i^r \quad (14)$$

where the u_i^{ta} and u_i^{to} are defined in the footnote³, $\mathcal{N}_i \setminus \mathcal{S}$ represents the set of all non-stakeholder neighbors of agent i , and u_i^r is given by Equation (5). Thus, the type-aware agent's attraction and orientation vectors are weighted between neighbors that are stakeholders and neighbors that are not stakeholders by the awareness parameter α_i .

The effect of an agent increasing α_i is to increase the influence of stakeholders, thereby amplifying human influence over the swarm. Because this amplification only happens when there are Q_i or more stakeholders in a type-aware agent's sensing neighborhood, the threshold Q_i acts as a nonlinear switch that improves the responsiveness of the swarm once enough agents are influenced by the human. We demonstrate empirically that this nonlinear switch based on the threshold Q_i improves the scalability of human influence over a swarm. Note that there is a trade-off between the responsiveness and the vulnerability of a swarm. A swarm that is highly responsive to changes in behavior made by only a few agents implies that if an adversary compromises a small percentage of the agents, or if a small percentage of the agents fail, then the swarm will be vulnerable to performing unwanted behaviors. We demonstrate that the embedded quorum response can act as a mechanism to limit the vulnerability of the swarm to agent failures.

5.2 Improved Scalability

To demonstrate improved scalability of human-swarm interactions, we compare how well the quorum sensing model (the model with $N - M$ type-aware agents) and the stakeholder model (the model with zero type-aware agents) allow switches between attractors. To investigate the scalability of these two models we repeated the experiment described in Section 4 using both the stakeholder model and the quorum sensing model for swarm sizes of $N = 100, 200, 300$, and 400 agents. We varied M from 10 to $N/2$ in 10 agent increments and varied ρ from between 0.1 and 1. Ten trials were performed for each M and ρ pair. When switching from torus to flock, we led the M stakeholders by orientation. When switching from flock to torus, we led the M stakeholders by attraction. To limit the number of parameters in these simulations we set $\alpha_i^{min} = 0$, $\alpha_i^{max} = 1$, and $Q_i = 0, \forall i$.

Figure 6 shows the minimum number of stakeholders necessary to switch group types and have 100% of the simulations remain in the new group type for $N = 100, 200, 300$, and 400. Results are shown for switching from a torus to a flock where M stakeholders are led by orientation (t2fo) and for switching from a flock to a torus where M stakeholders are led by attraction (f2ta). As the size of the swarm increases, the difference in scalability between the two models is very distinct. We see that switching between attractors using the quorum sensing model scales much better to larger swarm sizes. As the size N of the swarm increases, the number of agents that the human needs to interact with stays relatively constant for the quorum sensing model, but rapidly increases for the stakeholder model.

$${}^3u_i^{ta} = \alpha_i \frac{\sum_{j \in \mathcal{N}_i \cap \mathcal{S}} (c_j - c_i)}{\|\sum_{j \in \mathcal{N}_i \cap \mathcal{S}} (c_j - c_i)\|_2} + (1 - \alpha_i) \frac{\sum_{j \in \mathcal{N}_i \setminus \mathcal{S}} (c_j - c_i)}{\|\sum_{j \in \mathcal{N}_i \setminus \mathcal{S}} (c_j - c_i)\|_2}$$

$$u_i^{to} = \alpha_i \frac{v_i + \sum_{j \in \mathcal{N}_i \cap \mathcal{S}} v_j}{\|v_i + \sum_{j \in \mathcal{N}_i \cap \mathcal{S}} v_j\|_2} + (1 - \alpha_i) \frac{v_i + \sum_{j \in \mathcal{N}_i \setminus \mathcal{S}} v_j}{\|v_i + \sum_{j \in \mathcal{N}_i \setminus \mathcal{S}} v_j\|_2}$$

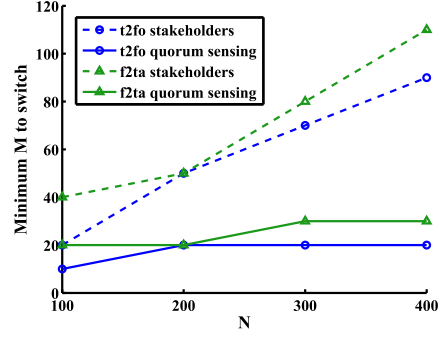


Figure 6: Minimum number of stakeholders, M , needed to switch from torus to flock where stakeholders are led by orientation (t2fo) and to switch from flock to torus where stakeholders are led by attraction (f2ta). Results compare the stakeholder model with the quorum sensing model.

5.3 Limited vulnerability


We now investigate whether using a quorum threshold allows us to limit the vulnerability of the swarm. We define limited vulnerability as the requirement that a swarm has an upper bound on the number of agents that can fail and not adversely affect the collective behavior.

To demonstrate that quorum sensing provides a way to limit the vulnerability of a swarm, we ran a series of experiments where M stakeholders attempt to switch the group type for different quorum thresholds Q . We experimented with values of Q between 0 and 6 and values of M between 0 and 50. Because of the large number of parameters in our model we restricted our analysis to parameter settings found to reliably allow the collective to switch between group types and remain switched. We used $\rho = 0.7$ when switching from a flock to a torus and $\rho = 0.5$ when switching from a torus to a flock. We led the agents by orientation when switching to a flock and by attraction when switching to a torus. We ran 10 replicates of each M and Q combination and used $\alpha_{max} = 1$ and $\alpha_{min} = 0$. We compare these results with the stakeholder model.


Figure 7 shows the probability of the swarm switching from a torus to a flock and from a flock to a torus. From this figure we see that the probability of switching decreases as Q increases. If we consider the stakeholders as malfunctioning or adversarial agents, then these results show that the quorum threshold Q provides a tunable parameter that controls the vulnerability/responsiveness of the swarm. By choosing a value for Q that is sufficiently high, we can limit the responsiveness of the swarm and prevent the swarm from switching group types unless there are sufficient agents making the change. Alternatively, if the goal is to have a swarm that is highly responsive, a low quorum threshold provides controlled switches with fewer agents.

5.4 Summary

We have shown that quorum sensing improves the scalability of human-swarm interactions. We have additionally shown that the quorum threshold Q provides a mechanism for limiting the vulnerability of a swarm. In light of our goal to find ways for a human to manage a bio-inspired robot



(a) Probability of switching from torus to flock



(b) Probability of switching from flock to torus

Figure 7: Switching between collective behaviors using only M stakeholders, or using M stakeholders and quorum sensing with $Q = 0-6$ ($N = 100$ agents).

swarm, these results indicate that it is possible for a human to select a switching probability before execution that balances scalability and vulnerability. Although this violates our goal of avoiding having a human broadcast parameters to the collective, it is useful since it allows a human to manage a trade-off prior to execution and still manage the flock and torus using a set of stakeholders during run-time.

6. RELATED WORK

Swarm models have been explored in a variety of fields and are typically capable of either flocking [26, 30, 23] or cyclic behavior [20, 19]. The model described in this paper is one of a small set of swarm models that exhibit multiple group behaviors [9, 29, 27]. This research concentrates on the flock and torus behaviors; however, there are other types of swarm behavior seen in the literature. Couzin’s model [9] has four group types: swarm, torus, dynamic parallel group, and highly parallel group. Couzin’s swarm behavior is analogous to a torus without orientation in our model, and the dynamic and highly parallel groups are simply two flavors of flocking. Strömbom [29] demonstrates that attractive forces between agents are sufficient to form swarms, flocks, and mills (torus-like formations where agents do not all rotate in the same direction). Strömbom also shows that adding a blind spot creates two additional group types: a torus and an interweaving chain-like structure. Romero et al. [27] present a swarm model that produces a swarm, a torus, a flock, and a flock that rotates around a stationary point. If we restrict our attention to models without centralized control,

global information, or explicit inter-agent communication, the group types mentioned above are the only group types we have found in the literature that emerge from swarm models capable of exhibiting multiple behaviors.

Couzin et al. [8] and Conradt et al. [5] explore leading a flock with a small number of informed agents. Couzin et al. show that their method of leading a flock scales well as group size increases, but do not consider leading a torus or changing group types. In the controls community much research focuses on consensus protocols [24] for flocking [23], but we have not found any research involving switching between attractors without communication or centralized control.

Some work has been done with communication-free flocking, but this work typically creates flocks that can be controlled by one or a few agents, which makes the flock vulnerable to leader failure. Gervasi and Prencipe [11] study distributed coordination and control without any communication or shared reference frame, but require that all agents can identify the leader agent and can be efficiently controlled by a single agent. Jadbabaie [15] provides mathematical results on the convergence of the Vicsek flocking model [30] to a single group direction as well as convergence conditions for the group to converge to a single leader’s direction.

Recent work has compared different methods for human-swarm interactions, but has either focused on flocking [25] or has relied on connectivity maintenance, rendezvous, and deployment algorithms that require a communication network between agents [18]. Nunnally et al. explore bandwidth constraints on swarm-to-human communications, but assume human communication with the entire swarm [22]. Our work extends and enhances the work by Kerman et al. [17], which provided preliminary results on human interaction with the flock and torus behaviors.

Several researchers have studied the effects of heterogeneity and quorum sensing on multi-agent interactions. Wu [3, 4] investigates the importance and role of heterogeneity in robot teams and shows that increased agent variation in task allocation problems can increase stability. Kumar [14] uses mathematical models of bio-inspired foraging tasks and quorum sensing to develop control strategies for redistributing agents among multiple sites, but does not investigate human influence over the agent behaviors.

7. CONCLUSIONS

We have presented a model of swarming that has two emergent behaviors: a flock and a torus. We also provided evidence that these behaviors are fundamental attractors of the swarm dynamics. Because these behaviors are attractors, a human operator can interact with the swarm by managing these attractors. We propose that human-swarm interactions should focus on managing higher level attractors of the swarm systems because it allows a human to abstract the details of individual agents and focus on managing the collective as a whole. We extended this work by presenting an application of quorum sensing to human-swarm interactions that increases the scalability of human-swarm interactions as well as provides a mechanism for allowing a human to balance a trade-off between vulnerability and responsiveness of the swarm to agent failures. Both the stakeholder and the quorum sensing models demonstrate the ability for a human to manage a swarm by managing its emergent behaviors. Future work should improve our static OoWiz model by examining dynamic interactions. In particular, user stud-

ies with humans will be necessary to fully understand the implications of controlling swarms by managing attractors.

8. ACKNOWLEDGMENTS

We appreciate the Office of Naval Research for funding this research under the Science of Autonomy Program. The findings, opinions, and conclusions of this paper do not necessarily reflect the views of the funding agency.

9. REFERENCES

- [1] N. Bode, D. Franks, and A. Wood. Limited interactions in flocks: relating model simulations to empirical data. *Journal of The Royal Society Interface*, 8(55):301–304, 2010.
- [2] D. S. Brown. Toward scalable human interaction with bio-inspired robot teams. Master’s thesis, Brigham Young University, Provo, UT, 2013.
- [3] A. Campbell, C. Riggs, and A. Wu. On the impact of variation on self-organizing systems. In *Fifth IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 119–128, 2011.
- [4] A. Campbell and A. Wu. Multi-agent role allocation: issues, approaches, and multiple perspectives. *Autonomous agents and multi-agent systems*, 22(2):317–355, 2011.
- [5] L. Conradt, J. Krause, I. Couzin, and T. Roper. Leading according to need in self-organizing groups. *The American Naturalist*, 173(3):304–312, 2009.
- [6] L. Conradt and T. Roper. Consensus decision making in animals. *Trends in Ecology & Evolution*, 20(8):449–456, 2005.
- [7] I. Couzin. Collective cognition in animal groups. *Trends in cognitive sciences*, 13(1):36–43, 2009.
- [8] I. Couzin, J. Krause, N. Franks, and S. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [9] I. Couzin, J. Krause, R. James, G. Ruxton, and N. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1–11, 2002.
- [10] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American J. of Mathematics*, vol. 79, no. 3, 1957.
- [11] V. Gervasi and G. Prencipe. Coordination without communication: The case of the flocking problem. *Discrete Applied Mathematics*, 144(3):324–344, 2004.
- [12] M. A. Goodrich, S. Kerman, and S. Y. Jung. On leadership and influence in human-swarm interaction. In *AAAI Fall Symposium Workshop on Human Control of Bio-Inspired Swarms*, 2012.
- [13] M. A. Goodrich, B. Pendleton, S. Kerman, and P. B. Sujit. What types of interactions to bio-inspired robot swarms and flocks afford a human? In *Proceedings of Robotics Science and Systems*, 2012.
- [14] M. Hsieh, Á. Halász, S. Berman, and V. Kumar. Biologically inspired redistribution of a swarm of robots among multiple sites. *Swarm Intelligence*, 2(2):121–141, 2008.
- [15] A. Jadbabaie, J. Lin, and A. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.
- [16] S. Kerman. Metrics for human interaction with bio-inspired swarms. Master’s thesis, Brigham Young University, Provo, UT, 2013.
- [17] S. Kerman, D. Brown, and M. Goodrich. Supporting human interaction with robust robot swarms. In *Proceedings of the International Symposium on Resilient Control Systems*, 2012.
- [18] A. Kolling, K. Sycara, S. Nunnally, and M. Lewis. Human swarm interaction: An experimental study of two types of interaction with foraging swarms. *Journal of Human-Robot Interaction*, 2(2):103–128, 2013.
- [19] H. Levine, W. Rappel, and I. Cohen. Self-organization in systems of self-propelled particles. *Physical Review E*, 63(1):017101, 2000.
- [20] J. Marshall, M. Broucke, and B. Francis. Formations of vehicles in cyclic pursuit. *IEEE Transactions on Automatic Control*, 49(11):1963–1974, 2004.
- [21] M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.
- [22] S. Nunnally, P. Walker, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara. Robotic swarm connectivity with human operation and bandwidth limitations. In *Proceedings of the AAAI Fall Symposium Workshop on Human Control of Bio-Inspired Swarms*, 2012.
- [23] R. Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.
- [24] R. Olfati-Saber, J. Fax, and R. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [25] B. Pendleton and M. Goodrich. Scalable human interaction with robotic swarms. In *To appear in the proceedings of the AIAA Infotech@Aerospace Conference*, 2013.
- [26] C. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH Computer Graphics*, volume 21, pages 25–34. ACM, 1987.
- [27] L. Romero, E. Forgoston, and I. Schwartz. Noise, bifurcations, and modeling of interacting particle systems. In *International Conference on Intelligent Robots and Systems*, pages 3905–3910. IEEE, 2011.
- [28] A. Steinfeld, O. C. Jenkins, and B. Scassellati. The “Oz of Wizard”: Simulating the human for interaction research. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, 2009.
- [29] D. Strömbom. Collective motion from local attraction. *Journal of theoretical biology*, 283(1):145–151, 2011.
- [30] T. Vicsek, *et al.* Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229, 1995.
- [31] P. Walker, S. Nunnally, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara. Investigating neglect benevolence and communication latency during human-swarm interaction. In *AAAI Fall Symposium Workshop on Human Control of Bio-Inspired Swarms*, 2012.