Robotic Swarms as Solids, Liquids and Gasses

Thomas B. Apker[*] and Mitchell A. Potter
US Naval Research Laboratory
Washington, DC 20375

Abstract

This paper presents an intuitive physics-inspired framework for controlling robot swarms that is based on the three major physical states of matter: solid, liquid and gas. An analogy is drawn between the states of matter and the basic swarming behaviors of clustering, translating and wandering. Mobility and localization requirements to achieve each of the states are specified.

Introduction

Physics-inspired, or physicomimetic, multi-robot control has two distinct advantages over animal-behavior inspired, or biomimetic swarm controls. First, it easier to engineer, in that the behaviors must be expressed in terms of Newtonian or Eulerian mechanics, which in turn allows the control designer to precisely specify the required level of mobility, localization and communication required for each agent during each phase of the mission. Secondly, based on our everyday experience with physical objects in our environment, the three major physical states of matter, solid, liquid and gas, represent a natural and intuitive means of describing the types of motions a swarm of mobile robots can perform as they cluster, transit or wander (Gage 1992).

This work presents a taxonomy of robotic swarm configurations and discusses the agent mobility and localization performance required to achieve them. In particular, it will demonstrate how to specify such biologically inspired behaviors as flocking and migrating in terms of physicomimetic interactions, and will describe how heterogenous teams of robotic agents can interact given vastly different capacities, such as ground robots coordinating with fixed wing air vehicles.

Background

A weakness of early physicomimetic approaches is that its basic agent model was the holonomic point particle (Spears and Gordon 1999). This lead to unrealistic performance expectations based on simulations (Spears et al. 2004; Wiegand et al. 2006), which tended to focus on equilibrium states of the swarm. This deficiency began to be systematically addressed by explicitly including a heading and turn-rate in the agent model, showing that limiting the agents' ability to turn reduced the quality of the lattice they would form using the same reactive controller as the holonomic particles (Ellis and Wiegand 2008). The authors (Apker and Potter 2011b) conducted a more thorough study of the impacts of motion constraints on swarming ability, and found it helpful to control agents' heading explicitly via virtual torques and a particle that "tows" the agent. In addition, they found that agent mobility has a strong impact on the types of formations available to the swarm (Apker et al. 2011).

Biomimetic approaches typically begin with motion-constrained agents and attempt to reverse engineer behavior models from observations of animals engaged in tasks such as hunting (Madden, Arkin, and MacNulty 2010), foraging (Liu and Hedrick 2011), or nest-site selection (Sasaki and Pratt 2011). There have been significant advances in developing algorithms that allow researchers to examine these behaviors in simulation (Luke et al. 2005), generally assuming noise-free estimates of the agents' own, neighbors' and targets' positions. However, the actual information flow into biological agents' in terms of the sensing, processing and communication required to produce these estimates is still a very active area of research (Jaroszewicz et al. 2007). Heterogenous agent interactions are quite rare and difficult to generalize to missions of human interest.

Most human interfaces that control live robots are based on artificial potential fields, such as Mission Lab (Ali and Arkin 2000), that employ a combination of agent-carried and map-fixed schema to direct agents to perform desired tasks and various abstractions to facilitate system design. A similar approach has been taken with physicomimetics to abstract low-level Newtonian parameters into concepts more intuitive to the operator to enable desired behaviors to emerge, such as forming an evenly spaced ring around a target (Kira and Potter 2009). Much of the work in physicomimetics involves finding proper distances between agents and relative heading angles, much like swarming and flocking algorithms from a biological perspective.

Physicomimetic Agents and Interactions

Like most swarm and robot team control approaches, physicomimetics is a guidance algorithm. As shown in figure 1,

[*]T. Apker is an NRC Postdoctoral Fellow at the US Naval Research Laboratory

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

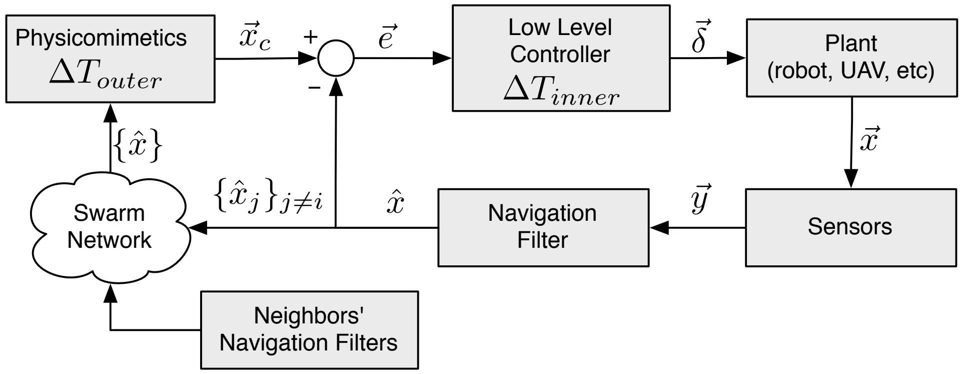

Figure 1: Block diagram of a typical physicomimetic implementation

its output is a commanded state, $\vec{x}_c$, that the inner-loop controller, ΔT_{inner}, attempts to meet by computing a set of actuator inputs, $\vec{\delta}$, to drive the difference between its estimated state, $\hat{x}$ and $\vec{x}_c$, (error signal $\vec{e}$) to zero. Sensors onboard the vehicle such as odometry, laser scanners and GPS produce measurement, $\vec{y}$, that are fused in the navigation filter to produce $\hat{x}$. With this in mind, we can use linear control theory to determine when $\vec{x}_c$ is in the reachable space of the inner loop within one pass through the outer loop, ΔT_{outer}.

Figure 1 succintly summaries the engineering challanges facing swarm control designers. Highly reactive and interactive swarms depend on accurate localization and timely message passaging between agents. Care must be taken to keep $\vec{x}_c$ in each agent's reachable space, otherwise the swarm as a whole will go unstable, especially in presense of disturbances (Li et al. 2011). The concept of local sensing, and possibly on-board estimation of neighbors' states, may reduce the dependance on electronic communication, but operator control assumes a swarm-spanning network and the problem of agents localizing their peers in real time remains a subject of active research.

Agent model

In (Apker et al. 2011), the authors described the "dumbbell agent," shown in figure 2 for physicomimetics control of real robots. Much like a biological agent, this version followed its "head," generally the front particle, towards its goal or equalibrium position. In some cases, linear force interactions, $\vec{F}$, were insuficient to generate or maintain swarm formations, and mission-level or localization, *e.g.* turning to point a camera at a landmark, require the vehicle to point in a specific direction. The dumbbell agent model accomplished this through a virtual torque, $\vec{T}$.

$$u_{j,k+1} = (u_{j,k} + \Delta t F_t/m)\mu_u; \quad (1)$$
$$\Omega_{j,k+1} = (\Omega_{j,k} + \Delta t \Sigma \vec{M}_j/I_{zz})\mu_\Omega \quad (2)$$

These forces, force couples and torques were combined to produce speed, u, and turnrate, Ω, commands as shown in equations 1 and 2. The forces on the front and back particle,

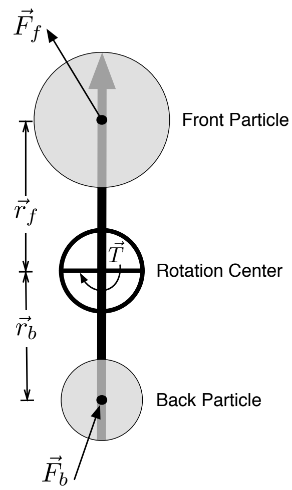

Figure 2: Diagram of the nonholomomic dumbbell agent

$\vec{F}_f$ and $\vec{F}_b$, respectively, were determined using the interaction models described below. The linear acceleration, or change in the speed command, Δu, is determined by the virtual mass m of the agent's particles, while the change in turnrate, $\Delta\Omega$, is a function of the torque; front and back force couples, themselves functions of the forces and distance from the rotation center to the particles, $\vec{r}_f$ and $\vec{r}_b$; and the virtual moment of inertia I. Further, we specified ranges of $u \in [u_{min}, u_{max}]$ and $\Omega \in [\Omega_{min}, \Omega_{max}]$ for each agent that were used both to bound $\vec{x}_c$ in the reachable space of the vehicle and to determine how heterogenous agents may interact.

Physicomimetic Interactions

The physicomimetics framework assumes that all trajectories are generated based on the interactions of particles connected to agents or salient features of the environment. This study focuses on three types of force laws specified in terms of solids, liquids and gasses. The physicomimetic force interactions considered here have three factors in common: (1) they are specified in terms of point singularities, (2) they are functions of the radial vector $\vec{r}$ between those points and (3) any individual interaction has a finite maximum range. The first two factors allow us to bound the complexity of each individual interaction, while the third enforces a sort of local measurement constraint that limits the overall computational cost of each agents' model. Torque or heading interactions are considered in terms of aligning agents, much like the Boids algorithm (Reynolds 1987), or facilitating attraction or repulsion.

The following sections consider three broad classes of interactions based on whether the objective is to cluster, transit or wander in the environment. Each of these cases places specific demands on the vehicle's inner loop, and so we will describe steps that can be taken at the interaction level to ensure that $\vec{x}_c$ remains reachable and each agent's error term $\vec{e}$ is minimized. Taking these steps results in slight losses in the quality of the swarm's lattice and smoothness of motion, but does result in predictions and control inputs that are more appropriate for swarm control.

Solid-Mode Interactions

The solid mode's primary interaction is the Lennard-Jones potential, which is a function of the distance between the agents, $\vec{r_{ij}}$, the equilibrium distance, σ, potential well depth factor, ϵ, and distance power, p, in equation 3 and plotted for several values of p in figure 3. The effectiveness of using the Lennard-Jones potential for large-scale swarm control with simulated watercraft using $p = 6$, also known as the six-twelve Lennard-Jones potential, has been previously demonstrated by (Frey et al. 2008). High values of p result in very narrow potential wells, which allowed the swarm to easily tesselate and avoid the second-order interactions that broader interaction ranges could cause.

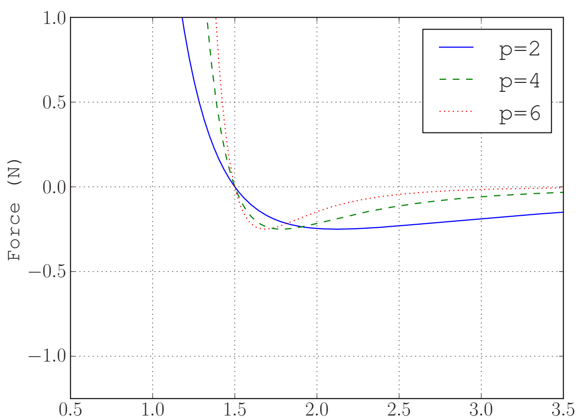

Figure 3: Lennard-Jones force with $p = 2, 4$ and 6.

$$\vec{F}_{ij} = \epsilon \left(\frac{\sigma^{2p}}{r_{ij}^{2p}} - \frac{\sigma^p}{r_{ij}^p} \right) \left(\frac{\vec{r}_{ij}}{r_{ij}} \right) \qquad (3)$$

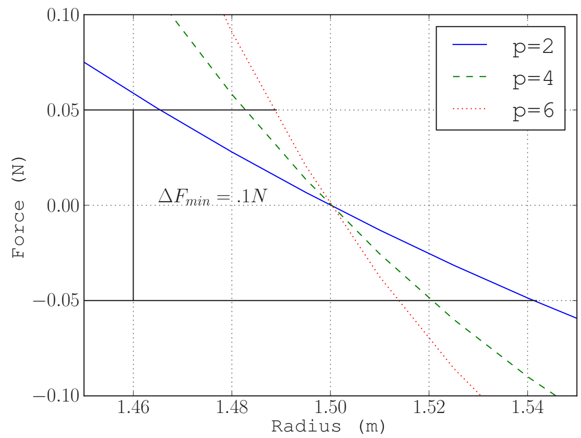

Figure 4: Deadband range of equation 3 for $p = 2, 4$ and 6.

However, narrow potential wells require very accurate localization of the interacting particles. Assuming the robotic agents use Kalman-type or particle filters for localization, there is a finite amount by which their estimates will shift with each uncertain measurement. Keeping the swarm at its equilibrium formation required us to define an explicit deadband distance Δx_{min} using equation 5 in which the agent can settle and adjust itself without disturbing its neighbors. To do this, we begin by finding the deadband in the force law about the equilibrium point, ΔF_{min}, defined in equation 4. Then, given information about the amount of noise we expect from our measurement system, we can find a suitable combination of p and ϵ for a given σ that maximizes both the quality of the swarm's formation and its stability in the face of noise and disturbances. An example of force law deadband is shown in figure 4.

$$\Delta F_{min} = \frac{m \Delta u_{min}}{\Delta T_{outer}} \qquad (4)$$

$$\Delta x_{min} = \frac{\Delta F_{min}}{\partial F / \partial x} \qquad (5)$$

From the operator's perspective, the goal of solid-mode interactions is to get the agents to cluster together with zero relative velocity, possibly in a specific formation. This implies that there is an overlap between all of the physical agents and interacting particles' range of speeds, *e.g.* fixed wing aircraft can cluster with other, similarly sized vehicles, but not ground vehicles or virtual particles at fixed points. In addition, the maneuverability of the cluster is significantly less than the agents. In the absense of explicit heading alignment, we found that the maximum pursuit speed of a swarm of five or more dumbbell agents in formation was less than 2% of the maximum agent speed, which for most mobile robots effectively means that the swarm must stop relative to an inertial frame, while holonomic agents experience a less dramatic loss in swarm speed with the number of agents. However, if the dumbbell agents explicitly align themselves using the Boids mutual-heading rule, the maximum linear pursuit speed is also the maximum agent speed. These results are summarized in figure 5.

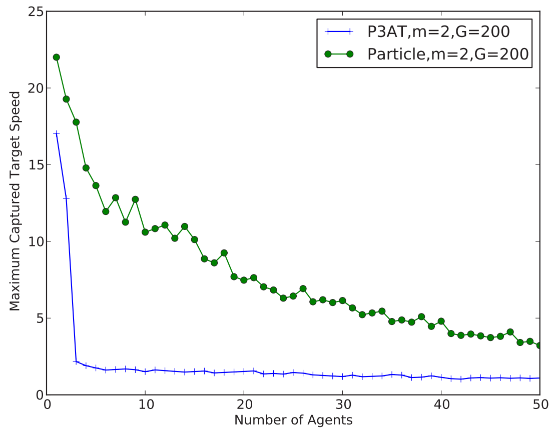

Figure 5: Movement rate of swarms of holonomic and dumbbell agents

The operators' commands in solid mode involve setting parameters for agent clustering. This allows behaviors such as gathering in an area for pickup, forming an antenna array or dispersal with equal spacing for sustained area coverage. As long as p and ϵ are valid for the values of σ chosen by the user or an algorithm such as the one presented in (Kira and Potter 2009), solid mode physicomimetics provides a simple and robust means of clustering a swarm of mobile robots.

Liquid Mode Interactions

One of the areas where biomimetic swarms perform well is in cases where the swarm needs to transit a large distance. Several algorithms for fish schooling and bird flocking produce robust paths for motion-constrained agents along curved paths and through cluttered environments. These paths and the agents following them exhibit a liquid-like quality, flowing around obstacles and rejoining as a cohesive body without explicit direction about which agents should pass each obstacle on its side. Physicomimetics generates similar trajectories for holonomic agents using radial force laws and separate particle types for obstacles and mobile agents.

Radial obstacle avoidance algorithms for nonholomic robots are generally more effective at preventing all motion by introducing local potential wells than encouraging translation around obstacles. To help in such circumstances, a "swirl force" with a radial component has been used to guide the agent around the obstacle instead of simple repulsion, for example to avoid a centrally located camera in (Martinson, Apker, and Bugajska 2011).

A substantial improvement to this for urban navigation has been developed by incorporating irrotational vortex panels into the obstacle model to generate streamlines, paths that do not intersect any obstacle or each other, in cluttered environments (Uzol, Yavrucuk, and Sezer-Uzol 2008). The fundamental obstacle particle in this model is the irrotational vortex, a tool used in aerodynamic modeling to approximate the behavior of flows over wing and tail surfaces.

Each vortex particle generates a tangential flow relative to the radius between itself and any point. Its direction in the global coordinate frame is defined as the cross product between $\vec{r}_{ij}$ and the unit vector normal to the plane $\hat{k}$. The strength of this interaction is a function of the vortex strength, Γ, divided by the distance between the j^{th} vortex and observation point, $\vec{x}_i$. In liquid mode physicomimetics, we treat this velocity component as a force input per equation 6. The strength of each vortex is set using the algorithm described in (Uzol, Yavrucuk, and Sezer-Uzol 2008) to guarantee that the sum of any additional fluid-type inputs, such as a uniform flow, source or sink, does not result in a force normal to the obstacle's mapped location.

$$\vec{F}_{ij} = \vec{r}_{ij} \times \hat{k} \left(\frac{\Gamma}{r_{ij}^2} \right) \tag{6}$$

The primary benefit of this liquid mode approach to translating in formation is that it is easier for an operator to explicitly direct the swarm to move in particular direction or to a particular point using the direction of a uniform flow and sink inputs. In cases where the streamlines curve too sharply for the agents to follow, they will simply move to adjoining streamlines and continue moving in the direction specified. The authors found in (Apker and Potter 2011a) that using liquid mode interactions with obstacles and solid mode interactions between fixed wing UAV-inspired agents produced an obstacle avoiding line-abreast formation that is well suited for plume detection and mapping.

Gas Mode Interactions

There are two cases when swarm agents cannot settle into desired locations or paths. The first is when the agents physically cannot stop relative to one another, such as a fixed wing UAV observing a stationary target or slow moving ground vehicle. The second occurs when the agents cannot provide a sufficiently accurate estimate of their own state to compute a meaningful $\vec{x}_c$. In these cases, we describe the swarm as being in gas-mode, as the most productive behaviors for the agents involve constant motion with as much change in direction as possible. In gas mode, we assume the controller provides heading inputs to the agents, while the inner loop controller or other local logic determines the best speed for the agent given hardware and environmental constraints.

Restricting the physicommietics controller to changing the heading while maintaining a constant speed aligns the swarm controller's commands with the reachable space of vehicles such as fixed wing UAVs. These have a preferred cruise speed for aerodynamic and stability reasons, and the linear behavior of their inner loop control is only valid for small deviations from it. In addition, there is great interest in developing rules to govern interactions between heterogenous systems such as air and ground vehicles that cannot settle into equilibrium positions or follow the same paths. In these cases, the constant and semi-random motion of gasses provides the only model for feasible interactions.

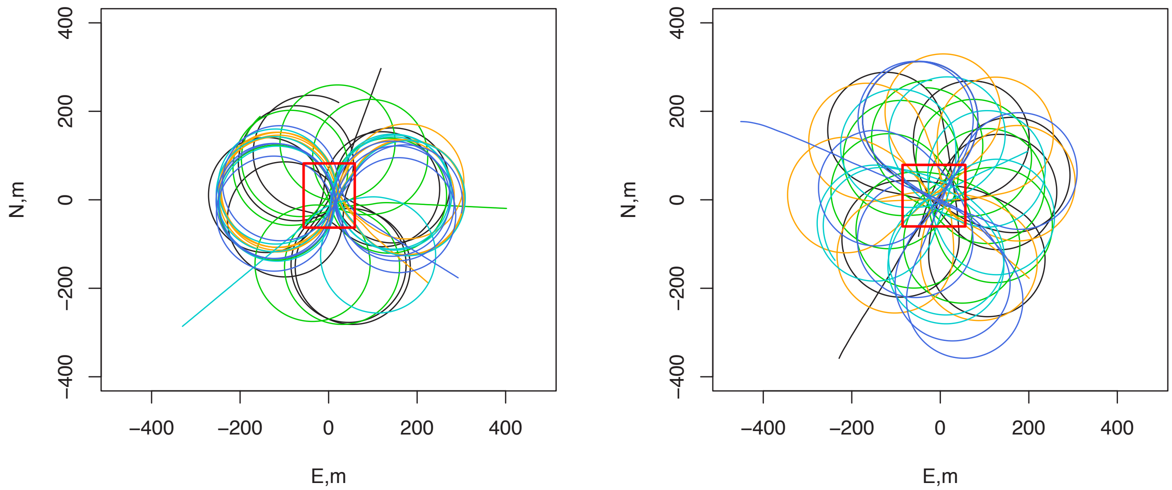

Figure 6: Plots of the paths flow by simulated UAV agents assuming small k_θ (left) and bang-bang control (right)

When agent localization allows it, such as small UAVs operating with GPS, we use the torsion-spring law described in (Apker et al. 2011). In this approach, agents are directed towards or away from a specific point or other agent by applying a virtual torque T that is a function of the angle θ between the vector between agents $\vec{r}_{ij}$ and its velocity $\vec{v}$ and a proportionality constant k_θ as shown in equation 8. For small values of k_θ, this approach drives agents into a spiral pattern in which they follow each other like beads on a string over fixed targets. At large values of k_θ, gas mode physicomimetics effectively becomes a bang-bang style controller in which agents turn as hard as possible away from each other or towards a target, producing a wider scattering of agents and thus more even area coverage as shown in figure 6. The choice of small vs. large k_θ depends on whether the user wants more time over a given target or broader coverage around the target.

$$\theta_{ij} = \cos^{-1}\left(\frac{\vec{r}_{ij}\cdot\vec{v}}{\|\vec{r}_{ij}\|\|\vec{v}\|}\right) \tag{7}$$

$$T = k_\theta \theta_{ij} \tag{8}$$

The heading-only gas mode is also helpful when localization is difficult and the goal is to allow agents to diffuse through a cluttered environment as quickly as possible. A compass or radio beacon, such as those used by the SenorFly system during their random walk deployment (Purohit et al. 2011), can provide a reasonable estimate of the desired heading, while onboard sensors provide information about obstacles and determine the most appropriate speed for the environment. Adding gas-mode interactions to the swarm "makes a virtue of necessity," and allows swarm designers to prepare for either the lack or loss of localization systems within the physicomimetics framework.

Conclusion and Future Work

The goal of this work was to provide a framework for robot team and swarm control that is species-agnostic and simple to explain to an operator with no background in robotics or biology. Based on everyday experience with solids, liquids and gasses, we presented a means of describing swarm behaviors that cluster, translate and wander along with the mobility and localization requirements to achieve them.

The dumbbell agent provided the basic mobility constraints of biological agents, and information from animal studies could be used to set the parameters of the interactions described above. Constraining agent motion resulted in physicomimetics lattice formations that more closely resembled the less-orderly clustering of biological agents than the even lattices of holonomic particles, suggesting that biological and physical inspiration lead to very similar results on robotic hardware. Defining translational behaviors in terms of simple fluid models allowed us to incorporate knowledge such as the map of an area into the swarm's behavior in a way that respected robot hardware constraints while still reproducing movements similar to biological flocking and schooling. Finally, we introduced a gas mode to handle cases where heterogenous agents could not operate in formation, wandering behaviors were desired or localization was not available to the agents.

In the future we plan to parameterize more biologically inspired behaviors in physicomimetics terms and define a means of tailoring interactions to specific vehicles. In particular, we are interested in using this solid-liquid-gas approach in conjunction with quorum sensing by cuing agents to transtion between modes based on their own sensor inputs.

Acknowledgments

This work was performed under Office of Naval Research Work Order N0001409WX30013.

References

Ali, K. S., and Arkin, R. C. 2000. Multiagent teleautonomous behavioral control. *Machine Intelligence and Robotic Control* 1(1):3–17.

Apker, T. B., and Potter, M. A. 2011a. An artificial physics approach to plume detection with fixed wing uavs. In *AIAA Guidance, Navigation, and Control Conference 08 - 11 August 2011, Portland, Oregon*.

Apker, T. B., and Potter, M. A. 2011b. *Swarm Intelligence Using Physicomimetics*. Springer. chapter Physicommimetics Motion Control of Physically Constrained Agents.

Apker, T. B.; Lennon, J.; Potter, M. A.; and Sofge, D. 2011. Past point models: Physicomimetics on nonholonomic vehicles. In *AIAA INFOTECH@Aerospace*.

Ellis, C., and Wiegand, R. P. 2008. Actuation constraints and artificial physics control. In *10th International Conference on Parallel Problem Solving from Nature (PPSN X)*, 389 – 398. Springer.

Frey, C. L.; Zarzhitsky, D.; Spears, W. M.; Spears, D. F.; Karlsson, C.; Ramos, B.; Hamann, J. C.; and Widder, E. A. 2008. A Physicomimetics Control Framework for Swarms of Autonomous Surface Vehicles. In *Oceans 2008, Vols 1-4*, 361–366. IEEE.

Gage, D. W. 1992. Command and control for many-robot systems. *Unmanned Systems* 10(4):28–34.

Jaroszewicz, A.; Sibilski, K.; Sibilska, A.; and Zyluk, A. 2007. Biomimic sensors guided flight stability and control for flapping wings autonomous micro air vehicle. In *AIAA Aerospace Sciences Meeting and Exhibit*.

Kira, Z., and Potter, M. A. 2009. Exerting human control over decentralized robot swarms. In *Proceedings of the 4th International Conference on Autonomous Robots and Agents*, 566–571.

Li, Z.; Vanness, J.; Caballero, D. E.; Hovakimyan, N.; and Stipanovi, D. 2011. Distributed target tracking and collision avoidance using multiple nonholonomic robots with uncertain dynamics. In *AIAA Guidance, Navigation, and Control Conference*.

Liu, S.-Y., and Hedrick, K. 2011. The application of domain of danger in autonomous agent team and its effect on exploration efficiencys. In *American Control Conference*, 4111–4116.

Luke, S.; Cioffi-Revilla, C.; Panait, L.; Sullivan, K.; and Bala, G. 2005. MASON: A multi-agent simulation environment. *Simulation: Transactions of the society for Modeling and Simulation International* 82(7):517–527.

Madden, J. D.; Arkin, R. C.; and MacNulty, D. R. 2010. Multi-robot System Based on Model of Wolf Hunting Behavior to Emulate Wolf and Elk Interactions. In *IEEE International Conference on Robotics and Biomimetics (ROBIO 2010)*, 1043–1050.

Martinson, E.; Apker, T.; and Bugajska, M. 2011. Optimizing a reconfigurable robotic microphone array. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 125–130.

Purohit, A.; Sun, Z.; Mokaya, F.; and Zhang, P. 2011. Sensorfly: Controlled-mobile sensing platform for indoor emergency response applications. In *10th International Conference on Information Processing in Sensor Networks (IPSN)*, volume 223–234.

Reynolds, C. 1987. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, 25–34. Association for Computing Machinery.

Sasaki, T., and Pratt, S. C. 2011. Emergence of group rationality from irrational individuals. *Behavioral Ecology* 22(2):276–281.

Spears, W. M., and Gordon, D. F. 1999. Using artificial physics to control agents. In *IEEE International Conference on Information, Intelligence, and Systems*, 281–288. IEEE.

Spears, W. M.; Spears, D. F.; Hamann, J. C.; and Heil, R. 2004. Distributed, physics-based control of swarms of vehicles. *Autonomous Robots* 17:137–162.

Uzol, O.; Yavrucuk, I.; and Sezer-Uzol, N. 2008. Collaborative target tracking for swarming MAVs using potential fields and panel methods. In *AIAA Guidance, Navigation and Control Conference and Exhibit*.

Wiegand, R. P.; Potter, M. A.; Sofge, D. A.; and Spears, W. M. 2006. A generalized graph-based method for engineering swarm solutions to multiagent problems. In *9th International Conference on Parallel Problem Solving from Nature (PPSN IX)*, volume 419 of *Lecture Notes in Computer Science*, 741–750. Springer.