Toward Human Interaction with Bio-Inspired Robot Teams

Michael A. Goodrich and Brian Pendleton
Computer Science Department
Brigham Young University
Provo, UT, USA
mike@cs.byu.edu

P. B. Sujit and José Pinto
University of Porto
Porto, Portugal

***Abstract*—In this paper, we formalize the problem of human interaction with bio-inspired robot teams (HuBIRT). The formalism applies to a large class of bio-inspired team dynamics and uses simple algebraic graph theory representations to distinguish between interagent influence, environmental influence, and operator influence. These representations lead to metrics for interagent cohesiveness and responsiveness to human input. We then select two different classes of team dynamics, physicomimetics which encodes dynamics using artificial physics, and a biomimetic structure which encodes dynamics using a model of fish behavior. We then demonstrate the relevance of the metrics by conducting a series experiments that demonstrate differences between leader and predator styles of human influence, and conclude with a comparison of nearest-neighbor topologies to metric-based topologies.**

***Index Terms*—human-robot interaction, biologically-inspired robotics**

I. Introduction

Two current research areas are receiving considerable attention in the recent literature: human-robot interaction (HRI) and bio-inspired robot teams (BIRT). HRI emphasizes the design of robot behaviors that support humans including when humans manage remote robots [1], [2].

BIRT research includes identifying principles and practices of biological *societies* [3], and then abstracting and encoding these principles in robots [4]. The resulting teams demonstrate so-called collective intelligence wherein simple robot behaviors produce colony-wide behaviors that appear collectively purposeful and goal-directed. Typical behaviors include swarming, flocking, foraging, and colony-building.

Steinberg has identified human-interaction with bio-inspired systems as an important research area of responsive, robust systems for complex surveillance and reconnaissance problems [5]. Research that combines elements of HRI with BIRT should allow humans to design robot teams that are both responsive and robust, yielding teams that can be efficiently managed by humans but that retain robust qualities in the presence of unreliability. Such research has been called cooperative robotics and human-swarm interaction, but we will call the work in this paper human-BIRT (HuBIRT) to emphasize human-centered BIRT design.

We will use an abstract *information foraging* task to evaluate HuBIRT performance. This information foraging task abstracts several reconnaissance and surveillance problems that are of current interest, including convoy protection [6] and contaminant tracking [7]. An information foraging problem is one where there are multiple tasks, represented as abstract resources, that appear at unknown locations in a spatial domain. Agents must discover the tasks and then assign a subset or subteam of the agents to perform the task. Each task takes time to complete, meaning that multiple agents must persist in the task for a satisfactory period of time before moving to another task. The resource size depletes at a rate based on the density of the number of assigned agents, that is, resource j depletes according to $S_j(t+1) = S_j(t) - \hat{N}s$ where $\hat{N}$ represents the number of agents within r_s meters of the resource location and $s > 0$ represents the amount of resource to be reduced per agent. For the experiments, we used $s = 0.001$ and $r_s = 5$. New tasks can appear anywhere in the domain at any time. Bio-inspired agents are capable of performing some aspects of this task by themselves, but are generally inefficient at the task without having some kind of human input.

The main contribution of this work is not a collection of mature algorithms and designs, but rather (a) formalizing the HuBIRT design problem for a large class of bio-inspired team models, (b) demonstrating the relevance of cohesiveness and responsiveness metrics, and (c) empirically studying cohesiveness and responsiveness properties for leader-based and predator-based human-BIRT interactions for both neighborhood-based and metric-based topologies. We perform experiments using two different types of BIRT teams: one physicomimetic and the other biomimetic.

II. Related Literature

There are several approaches to HuBIRT design. Centralized control of very large teams of agents has been called human-swarm interaction (HSI). Bashyal and Venayagamoorthy [8] presented an 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 is an example of practical HSI and uses swarm robotic technology to support firefighters using either proximate or remote operators [9].

Miller et al. advocate a playbook-style approach to human-robot teaming [10], 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) [11]. More complicated plays, like coordinated rendezvous or formation-following, have been applied to smaller teams [12]. The most obvious way to include a human in these approaches is for a human to control a leader agent who then influences other agents through conventional or bio-inspired means [13]. Managing patterns of behavior is another useful way to think about human interaction with bio-inspired teams. Recent work [13]–[16] explores centralized methods for influencing these patterns so that a human can guide the team to accomplish effectively some human-specified mission.

In addition to these centralized methods of HuBIRT design, there are approaches that take a more decentralized approach. For example, Barnes et al. explored how to modify potential fields in response to directions from a human operator [6]. Similarly, simulations have been performed of robots herding animals [17], [18], producing an ad hoc collection of observations, such as the fact that herding a flock using a single sheepdog requires very different behaviors than using a team.

III. Designing for Hu-BIRT

In this section, we present principles of bio-inspired interaction, formalize the principles, and identify some metrics associated with robust and responsive HuBIRT performance.

A. Bio-Inspired Principles

Sumpter identified several principles that describe biological systems which exhibit collective behavior [3]. In this section, we briefly describe and label a subset of Sumpter's principles. In the next section, we will use these brief descriptions to identify a discrete-time state space dynamical model.

Positive feedback is "imitation or recruitment behaviour [which] continues [until an] isolated behaviour is quickly subsumed by a mass of similar behaviors .. [yielding a set of] collective patterns."

Negative feedback is a counterbalance to positive feedback in that "positive feedback builds up a collective pattern ... [and] negative feedback ... stabilizes it. ... Negative feedback leads to ... stable output in the face of varied input."

Inhibition is negative feedback from other robots rather than the environment in that members "of a group exhibiting one type of behaviour can inhibit the behaviour of others."

Individual integrity means that each "... of the animals in a group is different, in terms of their ... previous experience [which can produce] ... different levels of response."

Response thresholds make allow agents to change "... behavior in response to a stimulus reaching some threshold."

B. Formalism

Sumpter's principles provide descriptive guidelines for what to look for in a system that exhibits collective intelligence, but they lack a formalism that could help HuBIRT design. Sumpter's principle of *integrity* suggests a state-space dynamic model where an agent's state is an encoding of previous "experience." Denote agent i's state at time t as x_t^i, and the vector of $N > 1$ agent states as $\mathbf{x}_t$. The precise encoding of state varies from model to model, but typical models include some notion of position and perhaps velocity [19]–[21].

Sumpter's principles of *positive feedback* and *inhibition* encode how the states of other agents affect agent i, and *negative feedback* and *response threshold* encode how external signals affect agent i. Thus, we adopt a state-space model that treats these two categories as two different functions. With some loss of generality, we will assume that the two categories are additive. Note (a) that this additive assumption is consistent with the physicomimentic and biomimetic models discussed below and (b) that this additive assumption is sufficient for swarming and flocking behaviors of homogeneous agents.

If we let u denote an external signal, then we can write the state dynamics model as

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t) + g(\mathbf{x}_t, u_t). \tag{1}$$

In terms of Sumpter's principles, Equation 1 represents positive feedback and inhibition in $f(\mathbf{x}_t)$, and negative feedback and response thresholds in $g(\mathbf{x}_t, u_t)$.

A key element of bio-inspired teams is that *collective intelligence emerges as a function of local interactions.* This manifests itself in a set of sparseness properties on f and g. Consider agent i and let $\neg i$ denote all agents other than agent i. Taking the i^{th} row from Equation 1 and partitioning $\mathbf{x}_t$ into x_t^i and $\mathbf{x}_t^{\neg i}$ yields,

$$x_{t+1}^i = f^i(x_t^i, \mathbf{x}_t^{\neg i}) + g^i(\mathbf{x}_t, u_t). \tag{2}$$

Since collective intelligence emerges from local interactions in BIRT, we will assume that the set of agents that affect agent i is small, that is, only a handful of agents in $\mathbf{x}_t^{\neg i}$ affect x_{t+1}^i. We call this assumption the *locality* assumption. Let A_t^i represent the adjacency matrix of the graph induced by $f^i(x_t^i, \mathbf{x}_t^{\neg i})$. We call A_t the *cohesiveness matrix* because it represents the interagent dynamics that determine the way collective intelligence of the team changes over time. Note that the performance of both artificial and biological systems depend critically on the structure of A_t [22], [23].

The vector u_t denotes all external signals, both those that come from the environment and those that may be specified by an operator. Thus, we separate u_t into two components: a subvector from the operator, u_t^{op}, and a subvector from the environment, u_t^{env}. In keeping with the biomimetic and physicomimetic models, we again assume that these components add. We call this assumption the *influence assumption* because it separates human influence from environment and inter-agent influence.

We next assume that $g^i(\mathbf{x}_t, u_t)$ is a function only of state x_t^i, that is, that external signals affect the next state of agent i only as a function of x_t^i and the external signals u_t. We call this the *autonomy assumption* because it says that agents can autonomously make decisions as a function only of their state and the signals that they receive independent of what other agents are doing. Thus, $g^i(\mathbf{x}_t, u_t) = g^i(x_t^i, u_t) = d_t^i(x_t^i, u_t^{\text{op}})+$

$e_t^i(x_t^i, u_t^{\text{env}})$. This paper will consider only operator and inter-agent influences; influences from the external environment will be treated in future work.

Let B_t denote the adjacency matrix induced by $\{d_t^i(x_t^i, u_t^{\text{op}}),\ i = 1, \ldots, N\}$. We call B_t the *management matrix* because it encodes the graph that represents which agents can be directly influenced by the actions of the operator. Note that B_t will often be sparse because decentralized control makes it unlikely that a human will be able to influence all agents simultaneously.

C. Performance, Cohesion, and Responsiveness

This paper takes an empirical rather than a theoretical approach to evaluating HuBIRT performance. Thus, we measure properties of (A_t, B_t) and correlate those properties to HuBIRT performance. One of the main results from decentralized control theory is that a fully connected group of agents often possesses strong collective properties like stability and cohesion [22]. Connectivity, in this context, refers to the connectivity encoded in the cohesiveness matrix and not in the management matrix.

Although connectivity in A_t appears to be necessary for cohesive behavior, it is not sufficient for HuBIRT. Cohesive and robust behavior must also be responsive to human input so that a human can direct collective behavior, which means that HuBIRT performance depends on both A_t **and** B_t.

According to Sumpter, **leadership** means that "key individuals ... catalyze and organize the group." In bio-inspired robot teams, two leadership models have been studied by others: lead-by-attraction and lead-by-repulsion. Lead-by-repulsion is more commonly referred to as *predation*. In this model, 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 models* and lead-by-attraction models *leader models*.

In our usage above, B_t captures the relationship between human input and the set of agents, and A_t captures interagent relationships. This raises the modeling question of whether we should (a) treat the leader/predator as a special type of agent whom the human influences through B_t and who then influences other agents through A_t, or (b) suppose that the human uses the leader/predator agent as an intermediary and then model how this intermediary influences the other agents through B_t, allowing A_t to encode interactions just between nominal agents. In this paper, we will adopt the latter because it allows us to distinguish between how the leader influences the agents and how that influence propagates throughout the collective; see Figure 1.

Many approaches to HuBIRT-related problems assume a communication channel that allows the human to broadcast and receive signals from all agents simultaneously. A more realistic model would allow a human to observe and communicate with only a subset of agents. In this model, the human is remote from the collective and the communication channel between human and the agents is bandlimited and may be subject to dropouts. We make two simplifying assumptions but note that future work should address these. First, we assume that local communication between agents is perfect. Second, we assume that the human can always observe the all agents' locations.

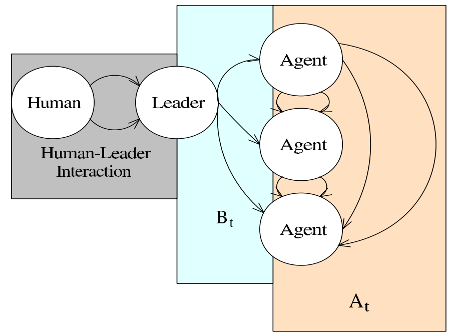

Fig. 1. Leader-mediated HuBIRT

As shown below, leader models have the desirable property that they tend to make the B_t matrix vary slowly in time; agents near the leader are attracted to and tend to stay close to the leader. This allows the leader to sustain influence over the collective.

By contrast, predator models tend to make the B_t matrix vary quickly in time; agents near the predator are repelled by the predator and try to get away from the predator's sphere of influence. Thus, as the predator moves, many agents escape from the sphere of influence, making it difficult for a predator to sustain influence over the collective.

D. Performance Metrics

We consider robustness, cohesion, and manageability.

Under HuBIRT, neighborhoods of any agent can be very dynamic. We are interested in how the influence of one agent is felt by other agents. Consider the time history of A_t over some temporal window,

$$\mathcal{A}_t = \sum_{k=0}^{T} A_{t+k}. \tag{3}$$

The i^{th} row of $\mathcal{A}_t(i)$ is a histogram of which of agent i's neighbors influence agent i during the time interval $[t, t+T]$.

$\mathcal{A}_t$ encodes a time-varying histogram of the cohesion matrix A_t. To support HuBIRT, a subset of agents should be influenced by the human, and this subset should exert sustained influence over their neighbors, and so on until the entire collective is influenced. Cohesiveness manifests itself in $\mathcal{A}_t$ in two ways: First, if all agents interact with most other agents almost all the time, then $\mathcal{A}_t$ will be "tall and uniform". Second, if each A_t is sparse, indicating just local interactions, then $\mathcal{A}_t$ should be "lumpy", meaning that agents tend to influence the same neighbors for multiple time steps.

Since $\mathcal{A}_t$ is only a function of the cohesiveness matrix, A_t, it doesn't explicitly represent the responsiveness of a bio-inspired team; rather, it represents the *potential* to be collectively responsive. In order for this potential to be fulfilled, the collective must also be responsive to the human's input. Since responsiveness is a function of B_t, let

$$\mathcal{B}_t = \sum_{k=0}^{T} B_{t+k} \tag{4}$$

denote the histogram of what agents the human is directly influencing. Sustained influence of a human over agents will be indicated by a "lumpy" or "pointy" visualization of $\mathcal{B}_t$.

IV. PHYSICOMIMETICS

Physicomimetics is a distributed control law that uses physics-based forces to control a swarm [24]. All agents are treated as point masses, all neighbors within the radius of C units influence behavior of agent i, and all agents outside of this radius have no influence. Each agent calculates the force acting on it by summing the forces from every other agent.

Since these agents are not goal-driven, efficient collective behavior requires human influence. We allow a human to influence agents by creating attracting or repelling forces on the agents. Human influence on any given agent depends only on the state of that agent and the human's influence, not on the state of any other agent, so the assumptions in Equation 2 hold.

A. Virtual Leader Management (VLM)

In VLM, the operator deploys a virtual agent that attracts all agents within a radius of attraction. Once the agents are attracted, the operator drags the virtual agent to the resource; once the resource is depleted the agents return to their nominal agent autonomy and re-distribute throughout the environment. Since the leader agent is virtual, the operator must continually broadcast to the agents to maintain influence.

B. Physical Leader Management (PLM)

In PLM, the operator can only communicate to a limited number of leader agents. A leader agent will recruit a number of agents and pull them to the resource location. The radius of attraction and the location of the resources is assigned to the leader by the operator. Leaders autonomously guide agents to resources; once a resource is consumed the agents redistribute.

The key difference between the PLM and VLM is the presence of delegation. In PLM, the human needs only to communicate remotely with the leaders through B_t, and all local interactions are propagated through A_t. By contrast, in VLM, the human needs to remotely communicate with all agents in its sphere of influence.

C. Virtual Predator Management (VPM)

VPM works similar to VLM but the virtual predator agent repels the agents inside its radius of influence. By repulsion, the virutal agent can push the agents towards the targets.

V. BIOMIMETICS

In this section, the scenario consists of 100 fish in a 120×120 area. Fish behavior is based on prioritized behavioral rules that tell a fish to change its desired direction as a function of the distance and direction of neighbors within a specified "zone of repulsion,", "zone of orientation," "zone of attraction", and "blind zone" [20]. Forces for agents are summative and include both attraction and repulsion components.

We consider leader- and predator-based human control, both of which are compatible with Equation 2.

A. Predator Management

The first control method involves using a single predator to steer groups of aligned fish. Fish are repelled by the predator if the predator is within a given radius. The predator moves slightly faster than the fish and can turn much more sharply.

Model parameters are set so that the fish are clustered in a small group, but if a predator gets close then they are repelled by this predator. Since the fishes' radius of mutual attraction exceeds the predator's radius of repulsion, the fish tend to stay close together even when the predator "chases" them.

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.

VI. RESULTS

In this section, we explore performance, cohesiveness, and management properties using simulations. Since the results do not include real robots, it is useful to briefly summarize the ecological validity of the experiments. The following assumptions apply to real robot systems: the robots are connected using a time-varying local topology with limited inter-agent communication, real human operators influence a small subset of agents, and additive dynamics allow for swarming and flocking behaviors. The following assumptions are not applicable to real robots: the human can observe the state of all agents, dynamics are noise-free, and data is obtained from only three human operators (all authors).

We begin by comparing the models in terms of task performance, measured as the time taken by a team to deplete the resource to zero, and of robustness, measured as the rate at which completion time increases as a function of the probability of communication loss or variations in resource distribution. We then explore characteristics of the cohesion and management matrices that correspond to good performance.

A. Cohesiveness and Manageability: Metric-Based Topologies

In this section, we explore how cohesiveness and manageability influence performance. We assume that agents are influenced by all other agents within a fixed distance of them, that is, the topology is metric-based.

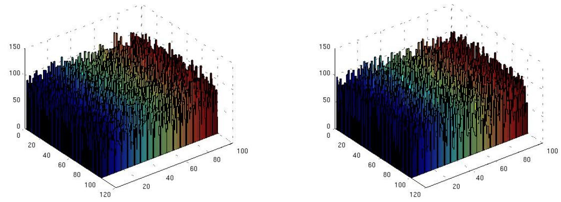

(a) (b)

Fig. 2. $\mathcal{A}$ for (a) leader model and (b) predator model

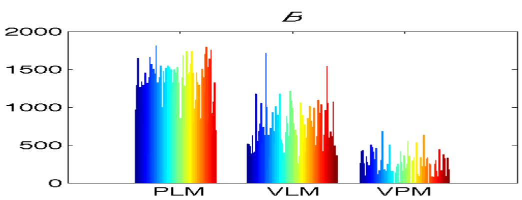

Fig. 3. Human influence through leader or virtual agent for different models

1) Cohesiveness: For the biomimetics model, we ran four experiments with each of the topologies for two minutes and gathered data. We recorded the A_t and B_t matricies at 1 second intervals. We performed a set of scenarios including an autonomous zig-zag leader and zig-zag predator through group. We also used a real human to lead the team.

We computed $\mathcal{A}$ as the sum of the A_t matrices over the entire two minutes of the simulation. As shown in Figure 2, the $\mathcal{A}$ matrix, which was computed over 120 time steps, is approximately 120 units tall. This indicates that the agents tended to stay close enough to each other that most agents were influenced by almost every other agent. This suggests strong cohesiveness. Moreover, the average values for the leader model are slightly higher than for the predator model, indicating that the predator model can cause "pockets" around a predator where agents are far enough apart that coherence decreases. Results from physicomimetics are similar.

2) Manageability: Manageability is reflected in how well humans maintain influence over agents. Under the physicomimetics model, we measured the $\mathcal{B}$ matrix by recording the number of time steps the human was able to influence the agents. Figure 3 shows that the agents using the PLM model have sustained human influence through the leader agents and hence their influence magnitude is higher compared to that of VPM model which has very low sustainable human influence through the virtual agent. The VLM model performs in between PLM and VPM models. The $\mathcal{B}$ matrices under biomimetics exhibit similar trends.

To reinforce how leader models sustain influence we computed a power spectral density for how the topology changed. For each agent, we created a time series of the number of new agents with whom the leader interacted plus the number of agents with whom the leader no longer interactedat each time index. We then computed the power spectral densities for these time series; see Figure 4. Note how the scales of

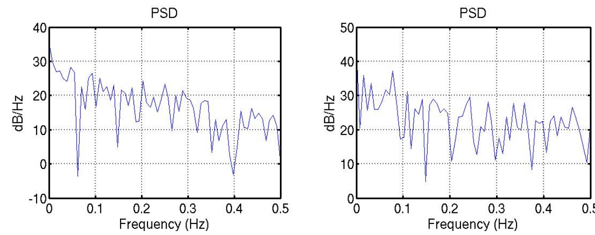

(a) Leader (b) Predator

Fig. 4. Power spectral density of $\mathcal{B}$ (biomimetic model) for (a) leader model and (b) predator model .

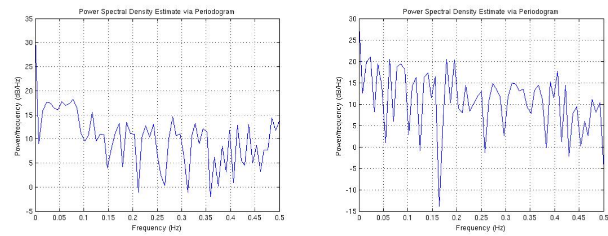

(a) Leader (b) Predator

Fig. 5. Power spectral density for nearest-neighbor topologies of $\mathcal{B}$ (biomimetic model) for (a) leader model and (b) predator model .

the two plots are different, indicating that the number and frequency of changes for predator models is much higher than for leader models. Simply put, predators cause agents to scramble, making it more difficult to sustain influence.

B. Cohesiveness and Manageability: Nearest-Neighbor Topologies

Ballerini noted that some flocks of birds appear to have local connections that are based on a handful of their nearest-neighbors rather than all birds within a fixed distance [23]. Thus, we repeated the above simulations but used nearest-neighbor topologies rather than metric-based topologies. The results for the bio-mimetic and physico-mimetic agents were similar, so we just present results for the former.

In comparing Figure 4 to Figure 5, note that the scale of the latter is much smaller. This indicates that neighborhoods change much less often for neighborhood-based topologies than metric-based. Simply put, neighborhood-based topologies are more cohesive, which is consistent with Ballerini's observation from nature [23].

C. Correlation with Performance

The temporal behavior of A_t and B_t indicate differences between leader-models and predator-models, and between metric-based and nearest-neighbor topologies. What is most interesting is that cohesive and manageable teams are more robust and responsive than less cohesive/manageable teams.

In the interest of space, we can only present a summary of performance and robustness results. Without fail, however, the performance of the HuBIRT teams for leader models was substantially better than predator models. Unsurprisingly in physicomimetics, physical leader models exhibited more robust

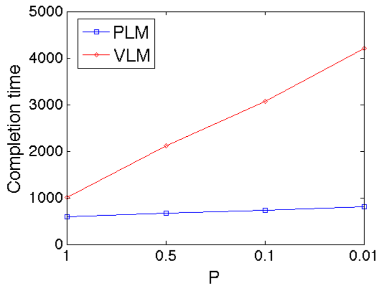

Fig. 6. Average response time for different models with varying communication probabilities.

performance than virtual leader models in the presence of communication failures, indicating that it is easier for a human to maintain influence over a single leader than over a large subset of agents in the presence of communication problems (see Figure 6). Moreover, subjective evaluations (note that this is only for three operators, all authors) strongly indicate that leader-based control of nearest neighbor topologies is substantially easier than predator-based control and marginally better than control of metric-based topologies.

VII. Conclusions and Future Work

This paper created a graph-based description for human interaction with bio-inspired robot teams (HuBIRT) based on biological principles. The model was encoded using two matrices, cohesiveness (A_t) and management (B_t), which were then used to characterize the performance of two different bio-inspired models. Cohesiveness is important because it allows agents to influence each other despite perturbations. Responsiveness means that agents can be managed effectively by a human via the management matrix. Using empirical studies, we provided evidence that a sufficient condition for HuBIRT to be responsive and robust is when (a) coherence is maintained through the connectivity of A_t and (b) influence is sustained through B_t; this may not be a necessary condition.

Future work should include exploring other communication models between the human and the agents, as well as among the agents themselves. Additionally, authority relations should be explored, perhaps by encoding interagent interactions as a directed graph. Issues of decoherence should also be considered, such as when it would be appropriate for a predator to split a collective into two subteams. Finally, multi-operator HuBIRT could be explored as a graph partitioning problem.

Acknowledgments

We appreciate ONR and ARL for funding. Results of this paper do not necessarily reflect the opinions of ONR and ARL.

References

[1] M. Cummings, S. Bruni, S. Mercier, and P. Mitchell, "Automation architecture for single operator, multiple UAV command and control," *The International C2 Journal*, vol. 1, no. 2, pp. 1–24, 2007.

[2] T. B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, 1992.

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

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

[5] M. L. Steinberg, "Biologically-inspired approaches for self-organization, adaptation, and collaboration of heterogeneous autonomous systems," in *Proceedings of SPIE Volume*, vol. 8062, April 2011.

[6] L. Barnes, M. Fields, and K. Valvanis, "Umanned ground vehicle swarm formation control using potential fields," in *Proceedings of the 15th Mediterranean Conference on Control & Automation*, Athens, Greece, July 2007.

[7] D. Zarzhitsky, D. F. Spears, and W. M. Spears, "Distributed robotics approach to chemical plume tracing," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005.

[8] 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.

[9] L. Alboul, J. Saez-Pons, and J. Penders, "Mixed human-robot team navigation in the GUARDIANS project," in *Proccedings of the International Conference on Safety, Security, and Rescue Robotics*, Sendei, Japan, October 2008.

[10] 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 *Proceedings of the 21st Digital Avionics Systems Conference*, vol. 2, November 2002, pp. 7E4–1 – 7E4–13.

[11] 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*, San Diego, March 2009.

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

[13] 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.

[14] Z. Kira and M. A. Potter, "Exerting human control over decentralized robot swarms," in *Proceedings of International Conference on Autonomous Robots and Agents*, 2009, pp. 566–571.

[15] J. Wang, M. Lewis, and P. Scerri, "Cooperating robots for search and rescue," in *Proceedings of AAMAS '06*, Hakodate, Japan, May 8-12 2005.

[16] 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 *Proceedings of AAAI Spring Symposium*, Stanford, CA, USA, 2006.

[17] R. Vaughn, N. Sumpter, J. Henderson, A. Frost, and S. Cameron, "Robot control of animal flocks," in *Proceedings of the International Symposium of Intelligent Control*, Gaithersburg, Maryland, 1998, pp. 277–282.

[18] J.-M. Lien, O. B. Bayazit, R. T. Sowell, S. Rodriguez, and N. M. Amato, "Shepherding behaviors," in *Proceedings of the IEEE International Conference on Robotics and Automation*, New Orleans, Louisana, 2004.

[19] L. Conradt, J. Krause, I. D. Couzin, and T. J. Roper, ""Leading According to Need" in self-organizing groups," *The American Naturalist*, vol. 173, no. 3, March 2009.

[20] 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*, vol. 218, no. 1, September 2002.

[21] N. W. F. Bode, D. W. Franks, and A. J. Wood, "Limited interactions in flocks: Relating model simulations to empirical data," *Journal of the Royal Society Interface*, vol. 8, 2011.

[22] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents part ii: Dynamic topology," in *Proceedings of the 42nd IEEE Conf. on Decision and Control*, Hawaii, December 2003.

[23] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic, "Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 105, no. 4, pp. 1232–1237, 2008.

[24] W. Spears, D. Spears, R. Heil, and W. Kerr, "An overview of physiocomimetics," in *Swarm Robotics*, ser. Lecture Notes in Computer Science, E. Sahin and W. Spears, Eds. Berlin: Springer, 2005, pp. 84–97.