2014 IEEE International Conference on Systems, Man, and Cybernetics
October 5-8, 2014, San Diego, CA, USA

Balancing Human and Inter-Agent Influences for Shared Control of Bio-Inspired Collectives

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

Shin-Young Jung and Michael A. Goodrich
Computer Science Department
Brigham Young University
Provo, UT USA
ldsrogan@gmail.com, mike@cs.byu.edu

***Abstract*—Human interaction with bio-inspired collectives provides an interesting setting for studying shared control. A human will often have knowledge of global objectives and high-level plans, but the collective will often have more detailed lower-level knowledge about the particulars of the situation at hand. Thus it is important to understand how control can be appropriately shared between the human and the collective. We analyze human interaction with bio-inspired collectives using graph theory, and propose that there are two human-side elements that determine how well control is shared: *span* and *persistence*. We additionally propose that there is a collective-side element that determines how well control is shared: *connectivity*. We study two examples of shared-control between a human and a bio-inspired collective: shaping a spatial formation and causing a collective to switch between stable collective states. Our empirical results show that span, persistence, and connectivity combine to affect (1) how influence is shared between the human and the collective and (2) the resulting success of human-collective interactions.**

I. Introduction

Human interaction with bio-inspired collectives provides an interesting setting for studying shared control. A human will often have knowledge of global objectives and high-level plans, but the collective will often have more detailed lower-level knowledge about the particulars of the situation at hand and will be robust and flexible in noisy, dynamic environments.

It is important to understand how control can be appropriately shared between the human and the collective. If a human has too much control, the collective loses its ability to react and change quickly and becomes a centralized system with a single point of failure. However, if the collective is given too much control, the human loses the ability to inject intelligence into the collective and higher-level mission objectives may suffer. Thus, shared control requires balancing human and agent influence in order to leverage the higher-level intelligence and ingenuity of a human operator as well as the collective intelligence and robustness of a bio-inspired collective.

We study two examples of shared-control between a human and a bio-inspired collective: shaping a spatial formation and causing a collective to switch between stable collective states. Shaping a spatial formation is useful when a human can provide small inputs to the collective's behavior to improve performance. An example would be an oil-spill recovery where a human positions a large number of cleaning robots to surround the spill and then adjusts the shape of the collective to maximize clean-up. This type of influence is fairly common; see, for example, [1]–[4].

The second way of influencing a collective, causing the collective to change from one behavioral state to a qualitatively different state, is useful when a human wants to have the collective switch between different tasks [5]–[7]. In the example of an oil-spill recovery, at some point the human will want the collective to stop their current behavior and move to a new location, either because the clean-up is done, or because there is a new task in a different location that requires immediate attention. This would require the entire collective to switch its behavioral state from "clean-up mode" to "travel mode".

The contributions of this work three-fold: (1) we provide a novel graph-theoretic partitioning for analyzing human-collective interactions and shared control, (2) we propose two graph-based factors that determine behavior, *span* and *connectivity*, and one temporal factor that determines behavior, *persistence*, and (3) we apply our graph-theoretic approach to two examples of shared-control between a human and a bio-inspired collective: shaping the spatial formation of a collective and causing a collective to switch between stable collective states. We show that span, persistence, and connectivity affect (1) how influence is shared between the human and the swarm and (2) the resulting success of the human-swarm interaction.

Note that related literature is reviewed throughout the paper rather than in a special section.

II. Measuring Influence

There are two human-side elements that determine how well control is shared: persistence and span. We define *persistence* as how long the human signals the swarm and *span* as how many agents the human influences. There is also one collective-side element that affects shared control: connectivity. We define *connectivity* in terms of the graph topology formed by agent interactions and influences.

A. An Influence-Based Partition

Key to analyzing shared control between a human and a collective is a way to measure (a) the balance between human-influence and inter-agent influence and (b) how these influences affect performance. To understand this, we follow [8] and use an influence-based partition of the adjacency matrix A

$$A(i,j) = \begin{cases} 1 & \text{if agent } i \text{ is influenced by agent } j \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

978-1-4799-3840-7/14/$31.00 ©2014 IEEE

We separate influence into three partitions: human-controlled, human-influenced, and collective-influenced. The *human-controlled* partition consists of the agents over which the human has direct control. The *human-influenced* partition consists of the agents in the collective that are either influenced indirectly by the human, through interactions with the human-controlled agents, or that receive information from, but are not directly controlled by, a human operator. Unlike, agents in the human-controlled partition, agents in the human-influenced partition may also be influenced by other agents in the collective. The last partition, *collective-influenced*, consists of the remainder of the collective: those agents that do not receive any information or control from the human and do not interact with the human-controlled agents. This partitioning yields

$$A = \left[\begin{array}{c|c|c} A_{hc} & 0 & 0 \\ \hline B_{hi} & A_{hi} & C_{hi} \\ \hline 0 & C_{ci} & A_{ci} \end{array}\right]. \tag{2}$$

The submatrix A_{hc} is the $N_{hc} \times N_{hc}$ adjacency matrix of the subgraph induced by the human-controlled agents. A_{hi} is the $N_{hi} \times N_{hi}$ adjacency matrix of the subgraph induced by the human-influenced agents; according to the partition, *span* equals N_{hi}. A_{ci} is the $N_{ci} \times N_{ci}$ adjacency matrix of the subgraph induced by the collective-influenced agents. B_{hi} is the $N_{hc} \times N_{hi}$ matrix encoding the influence of human-controlled agents on members of the collective, C_{hi} is the $N_{hi} \times N_{ci}$ matrix encoding the influence of the rest of the collective on human-influenced agents. Finally, C_{ci} is the $N_{ci} \times N_{hi}$ matrix that encodes the indirect influence of the human on the collective. We note that the zero blocks follow from our partitioning assumptions.

We investigate two classes of influence: the mediator class and the stakeholder class. The *mediator* class contains what prior work calls leaders, predators, and mediators [9]. In this class, a human directly controls certain agents, meaning that the A_{hc} partition is nonempty. The *stakeholder* class is defined by agents in the collective that a human can influence but not directly control, meaning that $A_{hc} = \emptyset$ and $B_{hi} = \emptyset$. These agents can also be influenced by their neighbors so $C_{hi} \neq 0$. We use the mediator class to shape a swarm's collective behavior and the stakeholder class to switch between collective states. We refer to agents that are not in the mediator or stakeholder class as nominal agents.

As defined above, this partitioning defines the *span* of human influence. We can also use this partitioning to evaluate the *connectivity* of the entire collective or any partition thereof as described in the following section. *Persistence* does not directly follow this partitioning, but will be shown to be relevant as we evaluate empirical studies.

B. Algebraic Connectivity

Given A, we can compute the graph Laplaican, $L = D - A$ where D is the diagonal degree matrix. The second smallest eigenvalue of L, which we denote by $\lambda_2(A)$, is called the algebraic connectivity [10]. The algebraic connectivity is often used to measure the connectivity of a graph and to provide a bound on the speed at which influence can flow through a graph [11]; however, this value is dependent on the number of agents in the partition. Chung [12] formulated a normalized version of the graph Laplacian such that the algebraic connectivity is always bounded between 0 and 2 for any size graph. This provides a nice metric for comparing results for different sizes and types of collectives because regardless of the number of agents in the collective or the number of agents in a particular partition we can compare the connectivity on a similar scale. We use both the Laplacian and the normalized Laplacian in this paper.

The normalized Laplacian is defined as

$$\mathcal{L} = D^{-1/2} L D^{-1/2} \tag{3}$$

where we use the convention that $D_{ii}^{-1} = 0$ if $D_{ii} = 0$. It can be shown that the second smallest eigenvalue of $\mathcal{L}$, λ_2, satisfies $0 \leq \lambda_2 \leq \frac{n}{n-1}$ where n is the number of nodes in the graph [12]. Thus, for large graphs, λ_2 is essentially bounded between 0 and 1.

Because the bound on λ_2 relies on having a symmetric[1] adjacency matrix, we compute the normalized algebraic connectivity using the symmetric adjacency matrix $\hat{A}$ where

$$\hat{A}(i,j) = \begin{cases} 1 & \text{if } A(i,j) = 1 \text{ or } A(j,i) = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Although we write algebraic connectivity as $\lambda_2(A)$, this is always computed from $\hat{A}$. While this is a simplification of the true network topology, we will show that this simplification still provides valuable information about the human and inter-agent influences that affect successful shared control.

C. Within- and Between-Partition Connectivity

Algebraic connectivity can be used to evaluate how agents influence each other within a partition or between partitions. *Within-partition influence* can be measured for any partition; for example, for nominal agents we use $\lambda_2(A_{ci})$. *Between-partition influence* can be measured across partitions; for example, influence between human-influenced agents and nominal agents can be measured using $\lambda_2(A_c)$ where

$$A_c = \begin{bmatrix} A_{hi} & C_{hi} \\ C_{ci} & A_{ci} \end{bmatrix}. \tag{5}$$

Algebraic connectivity is useful for measuring influence that is bi-directional, but to measure directed influence from the human to nominal agents we introduce complementary measures. For the mediator class of influence we measure indirect human influence as the fraction of the collective under the influence of human-controlled agents,

$$\mu_{hi} = \frac{N_{hi}}{N_{hi} + N_{ci}}. \tag{6}$$

This measures human influence over the swarm through B_{hi} and is useful for measuring human influence when $C_{hi} = 0$; this means that nominal agents are either totally influenced by the human-controlled agents or by their neighbors in the collective [13]. For the stakeholder class of influence, where

[1] Using this symmetric adjacency matrix is useful for evaluating connectivity of the team, but the asymmetries induced by the directed edges are important for avoiding the uncontrollability that can arise from purely symmetric graphs [8].

agents can be simultaneously influenced by both the human and the collective, we measure indirect human influence as

$$\mu_{ci} = \frac{1}{N_{ci}} \sum_{j=1}^{N_{ci}} \frac{N_{hi}^{j}}{N_{hi}^{j} + N_{ci}^{j}} \tag{7}$$

where $N_{hi}^{j} = \sum_{k=1}^{N_{hi}} C_{ci}(j,k)$ and $N_{ci}^{j} = \sum_{k=1}^{N_{ci}} A_{ci}(j,k)$. The value μ_{ci} tells us, on average, what proportion of a nominal agent's neighbors are human-influenced.

III. Shaping Collective Behaviors

Previous work has shown that mediators can be effective at changing the shape of the perimeter[2] of a collective [1], [15]. In this paper, we revisit the two types of mediator used in previous work: one that uses only repulsion (R-mediator) and a mediator that uses both repulsion and attraction (RA-mediator). All other agents use Couzin's dynamics with parameters set to yield a torus behavior [16][3].

We also consider a variation of nominal agents called S-agents [1]. The S-agents are influenced by an R-mediator, but deviate from Couzin-like agents by being able (a) to increase their speeds[4], and (b) to pay attention only to the nearest neighbor within their field of vision.

A. Human-influence

We measured the fraction of the collective that are being influenced by the mediator agents, μ_{hi}, for the R-mediator, the RA-mediator, and the S-agents. If a human is doing a good job controlling the mediators, then neither A_{hc} nor A_{hi} should be zero. Simulations were performed for 70 agents with a single mediator placed in the middle of a Couzin-like torus.

TABLE I. μ_{hi} FOR ONE MEDIATOR IN THE CENTER OF A COLLECTIVE.

Mediator Type	μ_{hi} Stationary	μ_{hi} Moving
R-Mediator	0.2445 (0.0474)	0.2361 (0.0528)
RA-Mediator	0.2585 (0.0484)	0.2527 (0.0514)
S-agents	**0.3192** (0.0560)	**0.0609** (0.0295)

The important results, in bold in Table I, indicate that when the mediator doesn't move, S-agents are more easily influenced than nominal agents. By contrast, when the mediator moves, the S-agents are less easily influenced than nominal agents. Subjective observations indicate that this is because S-agents pay slightly less attention to their neighbors than the nominal agents to each other. In other words, properly sharing control depends on both human influence and inter-agent interactions. In the next section we will show that because the S-agents have lower human influence, μ_{hi}, a low within-collective connectivity is required to enable shaping.

[2]Leaders and predators do not perform well from a shared control perspective. Leaders tend to turn all collectives into a swarm that does almost anything the human wants, and predators tend to fragment the collective [14], [15].

[3]Each nominal agent used the parameters $R^{rep} = 1$, $R^{ori} = 1$, $R^{att} = 20$, $s/unit = 4$, $\omega = 40°/sec$, $\theta = 180°$.

[4]In the simulations, speed was increased by 25% if the agent was more than $0.5 \cdot R^{att}$ from its nearest neighbor [1].

B. Connectivity

For the experiments in this subsection, we adopt a nearest neighbor topology [17] to manipulate within-partition influence among S-agents. Formally, each S-agent moves towards the centroid of its k-nearest neighbors and remembers the previous observed centroid of its neighbors in case it loses sight of them. The results in Figure 2 show that, as expected, $\lambda_2(A_{ci})$ increases as k increases, which means that the connectivity among the collectively influenced partition increases.

1) Expanding the collective: In the first experiment, 4 mediators were initially positioned at the center of a 40×40 unit area populated by 70 S-agents. The 4 mediators then moved to the corners of the square with speed one-fourth of S-agent speed; see Figure 1. As shown in Figure 2, $\lambda_2(A_{ci})$ increases

Fig. 1. Using mediators to expand $((a) \rightarrow (b) \rightarrow (c))$ or contract $((c) \rightarrow (b) \rightarrow (a))$ the collective shape.

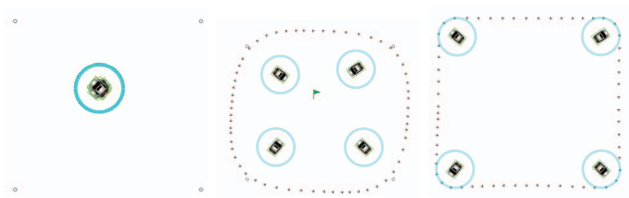

(a) (b) (c)

as with k. The constant values over time for $\lambda_2(A_{ci})$ simply indicate that the number of neighbors is fixed as long as there are neighbors within range. However, for $k \in \{5, 6, 7\}$ the connectivity begins to decrease at around 600 time steps. This corresponds to the shape breaking and the swarm fragmenting, indicating that shared control for this number of neighbors is not sustainable.

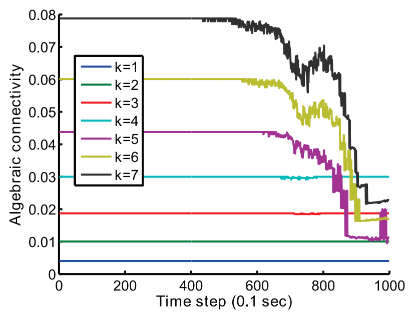

Fig. 2. Algebraic connectivity for different numbers of nearest-neighbors.

We repeated the experiment for $k \in \{5, 6, 7\}$ using a smaller mediator speed, and we were able to form a shape without fragmentation. However, we were not able to create a shape with more than 7 nearest neighbors even if we adjusted the speed of mediators.

To understand these results, note that, as illustrated in Figure 1, as the mediators move apart the torus stretches and begins to form a square. The stretching causes the S-agents, both those being influenced by a human and those not being so influenced, to move apart. The critical point in their movement occurs when there are not enough mediator-influenced S-agents to balance the collective influence among the remaining S-agents; simply put, the influence within the collective partition overcomes the influence between the human and the collective. This suggests that, by adding more mediators to increase span, it should be possible to sustain the expansion without fragmenting the collective.

Importantly, this result also tells us something about persistence. When mediators move more slowly, there is enough time for the shrinking number of nominal agents under the influence to percolate their motions throughout the collective. Moving more slowly makes the mediator's influence persistent enough to overcome forces within the collective partition.

2) Contracting the collective: We also tested contracting the collective by having the mediators start at the four corners and then move to the center. The results are shown in Table II.

TABLE II. TIME TO CONTRACT COLLECTIVE

# Neighbors	1	3	5	7
Time (s)	91.7	44.1	36.5	33.8

This effect is the reverse of what was observed in the previous section and can be simply explained by referring to Figure 3. As k increases from one to seven, agents respond to more neighbors causing a faster turning rate. The red dashed line describes the attraction vector towards each neighbor, and the green dashed line describes the calculated vectors. The solid green line is the final direction of movement.

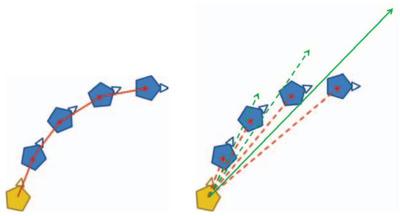

(a) nearest neighbor. (b) 4 nearest neighbors.

Fig. 3. Vector calculation shows that a 4-nearest neighbor topology results in higher curvature than a 1-nearest neighbor topology.

IV. SWITCHING BETWEEN COLLECTIVE BEHAVIORS

It is common for bio-inspired collectives to exhibit qualitatively different behaviors [6], [18]. Previous work [5], [19], [20] has investigated (a) a model of swarming that has two attractors, a flock and a torus, (see Figure 4), and (b) controlled switches between these attractors. The collective behaviors in this model result from members of the collective following simple attraction, orientation, and repulsion rules. This paper examines switching between a flock and torus from a shared control perspective.

A. Human interactions with stakeholders

In this section we examine shared control of a collective through stakeholders [5]. Stakeholders are not directly controlled by the human ($A_{hc} = \emptyset$ which implies $B_{hi} = \emptyset$) but rather are influenced by both the human ($A_{hi} \neq 0$) and by other agents ($C_{hi} \neq 0$). Each stakeholder receives a human supplied waypoint and has a persistence parameter $\rho \in [0, 1]$ that determines the priority of moving towards the waypoint versus reacting to nearby agents. If ρ is high, then a stakeholder responds more to human commands. If ρ is low, then a stakeholder responds more to its neighbors.

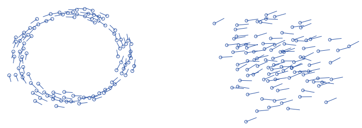

(a) Torus (b) Flock

Fig. 4. The two behaviors formed by the swarm model in [19]. Agent headings are represented by straight lines.

Since $A_{hc} = \emptyset$ and $A_{hi} \neq 0$, our analysis focuses on the connectivity of the human-influenced partition, the collective-influenced partition, and the entire collective. We measure indirect influence from human to nominal agents using μ_{ci}.

B. Switching from Flock to Torus

Using a type of stakeholder designed for the task, Brown et al. [5] used stakeholders to switch from a flock to a torus. Figure 5 shows the percentage of trials that caused a switch from a flock to a torus when N_{hi} randomly chosen stakeholders were under the influence of a human, and the percentage that remained a torus after human input was removed. The (x, y) plane has axes of span, N_{hi}, and persistence, ρ. Results are for a collective of 100 agents.

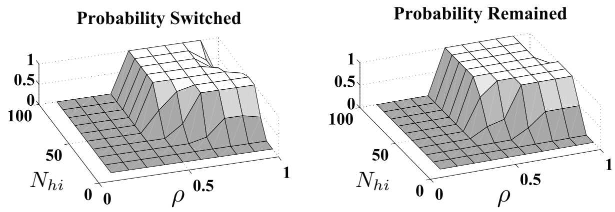

Fig. 5. Switching from flock to torus [5].

We selected four pairs of values for ρ and N_{hi} that illustrate when shared-control works and when it breaks-down when switching from a flock to a torus: (1) $N_{hi} = 30$ and $\rho = 0.3$, insufficient priority put on human influence; (2) $N_{hi} = 90$ and $\rho = 1$, insufficient priority put on inter-agent influences and too much span; (3) $N_{hi} = 40$ and $\rho = 0.7$, good shared-control; and (4) $N_{hi} = 20$ and $\rho = 1$, insufficient priority put on inter-agent influences and not enough span.

Figure 6 shows μ_{ci}, $\lambda_2(A_c)$, $\lambda_2(A_{hi})$, and $\lambda_2(A_{ci})$ for the above parameters. For comparison, horizontal dash-dotted lines show the average connectivity for no human influence ($N_{hi} = 0$) for a flock ($\lambda_2(A_{ci}) \approx 0.7035$) and torus ($\lambda_2(A_{ci}) \approx 0.4393$). Vertical dashed lines ($t = 250$ and $t = 2250$) show when human control is added and removed.

Examining the influence partition gives us insight into the dynamics of the collective. First, Figure 6(a) shows that when $N_{hi} = 30$ and $\rho = 0.3$ the shared control is unbalanced in favor of the collective and the human is unable to influence a switch

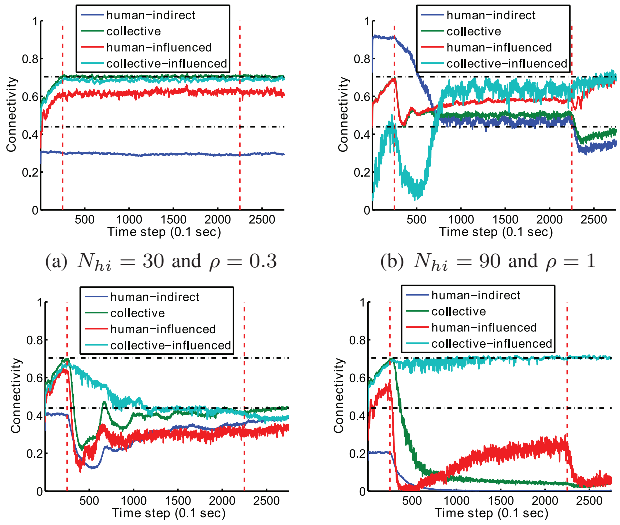

(a) $N_{hi} = 30$ and $\rho = 0.3$ (b) $N_{hi} = 90$ and $\rho = 1$

(c) $N_{hi} = 40$ and $\rho = 0.7$ (d) $N_{hi} = 20$ and $\rho = 1$

Fig. 6. Connectivity for switching from a flock to a torus.

from a flock to a torus. Second, Figure 6(d) shows that when $N_{hi} = 20$ and $\rho = 1$, there is distinct loss of connectivity. The dip in connectivity for the human-influenced partition after $t = 2250$ is because the stakeholders are no longer kept connected by the human influence.

Third, in Figure 6(b) we see that when $N_{hi} = 90$ and $\rho = 1$ the collective connectivity looks like a torus but the nominal agent and stakeholder connectivity values still resemble a flock. Figure 7 shows the actual behavior of the collective during and after human influence. We see that the group is almost torus like, but unstable. When the human influence is released, the collective turns back into a torus. This shows the value of looking at the connectivity of partitions and of the entire collective to understand how control is shared and why performance might suffer.

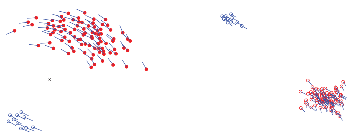

(a) While under human influence (b) After human influence released

Fig. 7. Switching from a flock to a torus with 90 stakeholders (solid red) with $\rho = 1$.

Fourth, results in Figure 6(c) for $N_{hi} = 40$ and $\rho = 0.7$ show that all connectivity measures converge toward the torus connectivity indicating good shared control. This convergence is key to understanding shared control. It shows that for a properly balanced collective, there is a high correlation between the indirect human influence measured by μ_{ci} and the human-influenced stakeholder connectivity. The correlation is also high for the entire collective, but the correlation has a time lag. The time lag is caused by the amount of time that the nominal agents require to respond. The time lag is the *persistency* requirement for controlling the collective—the human must persist long enough for the entire collective to converge to the new state.

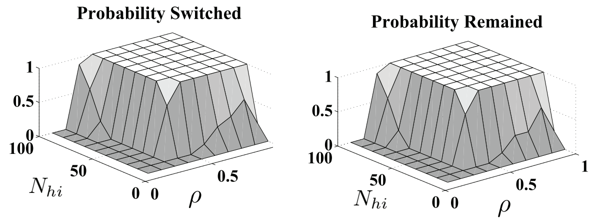

Fig. 8. Switching from torus to flock [5].

C. Switching from Torus to Flock

Properly tuned stakeholders can also be used to turn a torus to a flock [5]. Figure 8 shows the percentage of trials that caused a switch from a torus to a flock when N_{hi} randomly chosen stakeholders were under the influence of a human, and the percentage that remained a flock after human input was removed. As before, the (x, y) plane has axes of span, N_{hi}, and persistence, ρ.

We again selected values for ρ and N_{hi} that illustrate when shared-control works and when it breaks-down when switching from a torus to a flock: (1) $N_{hi} = 30$ and $\rho = 0.3$, not enough span or human-influence; (2) $N_{hi} = 50$ and $\rho = 0.5$, a good balance between human-influence and inter-agent influence; (3) $N_{hi} = 10$ and $\rho = 0.6$, too much human influence and not enough span; and (4) $N_{hi} = 40$ and $\rho = 1$, too much human-influence and not enough inter-agent influence.

The partition-based connectivity results for switching from a torus to a flock are shown in Figure 9. First, when $N_{hi} = 30$ and $\rho = 0.3$ the inter-agent connectivity is too strong and the stakeholders cannot cause a switch to a flock. Second, when $N_{hi} = 50$ and $\rho = 0.5$ there is a transition between torus and flock, and convergence of connectivity parameters indicating good shared control. Third, when $N_{hi} = 10$ and $\rho = 0.5$ the span is too low and the human-influence is strong resulting in fragmentation of the collective.

When $N_{hi} = 40$ and $\rho = 1$ the collective connectivity drops below the connectivity of a normal torus, then gradually increases, and then drops again when the human influence is released. Figure 10 illustrates why this happens. The high persistence parameter, $\rho = 1$, causes the stakeholders to ignore their neighbors and never fully form a flock. Thus, the stakeholder connectivity is low despite the nominal agents clumping and forming a stable flock behind the stakeholders. Once the human influence is released the stakeholders in front are attracted to and orient to their neighbors behind them, forming a large clump of nominal agents followed by a distant clump of stakeholders.

V. Conclusions and Future Work

Using two examples of human interaction with bio-inspired collectives, we demonstrated the importance of span, persistence, and connectivity in achieving balanced, successful shared control. We showed that the span of human control must be high but not too high. The mediator results in Section III-B1

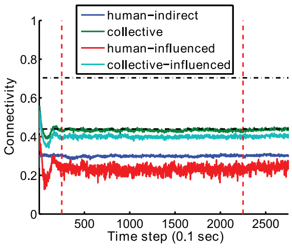

(a) $N_{hi} = 30$ and $\rho = 0.3$

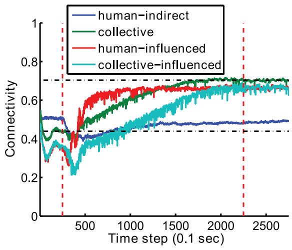

(b) $N_{hi} = 50$ and $\rho = 0.5$

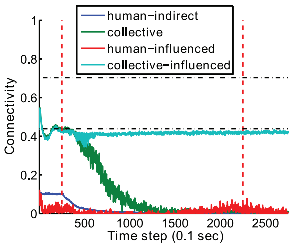

(c) $N_{hi} = 10$ and $\rho = 0.6$

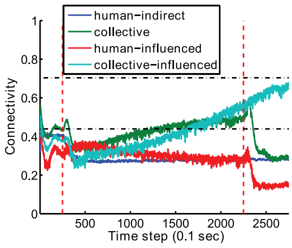

(d) $N_{hi} = 40$ and $\rho = 1$

Fig. 9. Connectivity for switching from a torus to a flock.

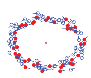

(a)

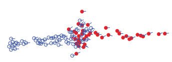

(b)

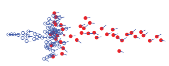

(c)

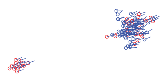

(d)

Fig. 10. Torus to flock with $N_{hi} = 40$ and $\rho = 1$: (a) Human influence applied to torus, (b) Torus switches to elongated flock, (c) The flock never fully forms, (d) Human influence is released.

demonstrated that low span can cause the perimeter of the collective to fragment. When using stakeholders, small spans resulted in either fragmentation or in the collective ignoring the human influence. However, a large span also presents problems (see Figure 7) because it disrupts the natural behavior of the collective and causes unstable behaviors.

We also showed that the persistence of human interactions must be high, but that over aggressiveness leads to failure. We found that decreasing the speed of mediators made the human influence persistent enough to continue to be able to shape the collective. All interactions with stakeholders required persistence to allow the collective to converge to a new state; however, we showed that high persistence with a small span resulted in fragmentation.

Finally, we showed that between-partition and within-partition connectivity must be balanced. We showed that because S-agents have low human influence when under the influence of a moving mediator, the within-collective connectivity must also be low in order to successfully shape the collective. Stable switches between collective states resulted when between-partition and within-partition connectivities were balanced and converged over time (see figures 6(c) and 9(b)).

REFERENCES

[1] S.-Y. Jung, D. S. Brown, and M. A. Goodrich, "Shaping couzin-like torus swarms through coordinated mediation," in *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, 2013, pp. 1834–1839.

[2] J. Wang and M. Lewis, "Assessing coordination overhead in control of robot teams," in *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 2007, pp. 2645–2649.

[3] 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.

[4] 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.

[5] D. S. Brown, S. C. Kerman, and M. A. Goodrich, "Human-swarm interactions based on managing attractors," in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. ACM, 2014, pp. 90–97.

[6] 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.

[7] R. P. Wiegand, M. A. Potter, D. A. Sofge, and W. M. Spears, "A generalized graph-based method for engineering swarm solutions to multiagent problems," in *Lecture Notes in Computer Science*. SpringerLink, 2006, vol. 4193.

[8] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt, "Controllability of multi-agent systems from a graph-theoretic perspective," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 162–186, 2009.

[9] M. A. Goodrich, S. Kerman, and S.-Y. Jung, "On leadership and influence in human-swarm interaction," in *2012 AAAI Fall Symposium Series*, 2012.

[10] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak Mathematical Journal*, vol. 23, no. 2, pp. 298–305, 1973.

[11] M. Mesbahi and M. Egerstedt, *Graph theoretic methods in multiagent networks*. Princeton University Press, 2010.

[12] F. R. Chung, *Spectral graph theory*. American Mathematical Soc., 1997, vol. 92.

[13] M. A. Goodrich, B. Pendleton, P. Sujit, and J. Pinto, "Toward human interaction with bio-inspired robot teams," in *Systems, Man, and Cybernetics (SMC), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2859–2864.

[14] 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*, Sydney, Australia, June 2012.

[15] S.-Y. Jung and M. A. Goodrich, "Multi-robot perimeter-shaping through mediator-based swarm control," in *Advanced Robotics (ICAR), 2013 16th International Conference on*. IEEE, 2013, pp. 1–6.

[16] I. Couzin, J. Krause, R. James, G. Ruxton, and N. Franks, "Collective memory and spatial sorting in animal groups," *Journal of Theoretical Biology*, vol. 218, no. 1, pp. 1–11, 2002.

[17] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini *et al.*, "Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study," *Proc. of the National Academy of Sciences*, vol. 105, no. 4, p. 1232, 2008.

[18] T. D. Seeley, *Honeybee democracy*. Princeton University Press, 2010.

[19] 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.

[20] D. S. Brown and M. A. Goodrich, "Limited bandwidth recognition of collective behaviors in bio-inspired swarms," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multiagent Systems*. ACM, 2014, pp. 90–97.