

Multi-Agent System Dynamics: Bifurcation and Behavior of Animal Groups [★]

Naomi Ehrich Leonard

*Mechanical and Aerospace Engineering, Princeton University,
Princeton, NJ 08544 USA (Tel: +1 609-258-5129; e-mail:
naomi@princeton.edu)*

Abstract: Systematic design of decentralized feedback for coordinated control of multi-agent systems has much to gain from the rigorous examination of the nonlinear dynamics of collective animal behavior. Animals in groups, from bird flocks to fish schools, employ decentralized strategies and have limitations on sensing, computation, and actuation. Yet, at the level of the group, they are known to manage a variety of challenging tasks quickly, accurately, robustly and adaptively in an uncertain and changing environment. In this paper we review recent work on models and methods for studying the mechanisms of collective migration and collective decision-making in high-performing animal groups. Through bifurcation analyses we prove systematically how behavior depends on parameters that model the system and the environment. These connections lay the foundations for proving systematic control design methodologies that endow engineered multi-agent systems with the remarkable features of animal group dynamics.

1. INTRODUCTION

Multi-agent system dynamics have gained significant attention in the control community in recent years. This is due to the rich theory associated with decentralized feedback control and system performance as well as to the growing number of important and challenging applications in cooperative control of networked dynamical systems, from robotic vehicle networks to electric power networks to synthetic biological networks (Antsaklis and Baillieul [2004, 2007], Bullo et al. [2009], Mesbahi and Egerstedt [2010], Zachary et al. [2011]). It is common in multi-agent system dynamics to limit each individual agent in the system in terms of what it can decide on its own, what it can do on its own, and what it can measure on its own about its local environment. A key objective is then to prove that through judicious design of feedback and interaction among the agents, the decentralized networked multi-agent system can function at a very high level, meeting demanding performance criteria in complex tasks.

The present paper reviews a selection of recent work by the author, with collaborators from engineering and biology, to develop abstract models and methods for rigorous examination of mechanisms of multi-agent dynamics. The focus is on high-performing dynamics drawn from observations and understanding of collective behavior in animal groups. The study of collective animal behavior is particularly relevant to decentralized feedback control design: animals typically apply decentralized strategies and have limitations on their ability as individuals to compute, act, and sense, and yet as a group they adapt their behavior with accuracy, speed and seemingly little effort.

A central goal of the work presented is to derive analytically tractable dynamic models that capture the mechanisms of collective animal behavior and lend themselves to systematic study of sensitivity of performance to critical parameters of the system and the environment. This creates foundations for developing provable and systematic control design methodologies that enable engineered multi-agent systems to inherit some of the remarkable features of animal group dynamics.

The dynamics of collective animal behavior are typically nonlinear due to nonlinearities in individual dynamics, nonlinearities in interaction dynamics, nonlinear coupling between the individual dynamics and the interaction dynamics, and nonlinearities in the configuration space. Models also tend to be high-dimensional especially when the number of agents in the group is large.

Approaches such as time-scale separation, or mean-field approximation, can be used to derive reduced-order models. Bifurcation theory can then be applied to isolate the influence of system and environmental parameters on structural changes in the collective behavior. When bifurcation parameters, such as the feedback gains of individual agents, are endowed with their own dynamics, either by hypothesis to model animal groups or by design to control engineered networks, the resulting system dynamics describe adaptive collective behavior, where conditions for high performance can be proved in a possibly changing environment. Analytic approaches to investigating adaptive behavior in biology, e.g., evolutionary dynamics and analysis of evolutionarily stable solutions, can be integrated with control theoretic approaches to advance a theory of multi-agent system dynamics.

Animal groups provide enormous inspiration when viewed as high-performing multi-agent systems implementing decentralized feedback control in a changing and uncertain environment. From bird flocks to fish schools, animal

[★] Plenary paper, IFAC Symposium on Nonlinear Control Systems (NOLCOS), Toulouse, France, September, 2013.

This work was supported in part by ONR grant N00014-09-1-1074, ARO grant W911NG-11-1-0385, and NSF grant ECCS-1135724.

groups exhibit a superior ability to manage a variety of challenging tasks, from foraging to migration to predator evasion, that individual animals would not be able to manage on their own (Krause and Ruxton [2002]). Observations of collective behaviors suggest that animal groups perform exceptionally well in the trade-off between speed and accuracy and between robustness to disturbance and responsiveness to environmental change (e.g., Ballerini et al. [2008]). Biologists attribute the remarkable behavior of these groups in large part to the social interactions among individuals. Indeed, interactions among many relatively limited individuals have been shown to yield collective dynamics that are much more than the sum of their parts (Parrish and Edelstein-Keshet [1999]).

The work reviewed in the present paper focuses on dynamics of collective migration and collective decision-making. Many species of birds, fish, invertebrates and mammals rely on social interactions during migration, especially when it is costly for individuals to measure stochastic environmental cues such as nutrient and thermal gradients, magnetic fields, odor and visual markers (Fryxell and Sinclair [1988], Holland et al. [2006]). Migration can be an adaptive response to seasonal variability, competition for dynamically changing resources, and shifting focus on habitats for breeding, where animals must learn, agree on and follow a single migration route. Collective migration has been shown to be evolutionarily stable for certain environmental conditions in the case that only a subset of individuals invest in measuring the costly environmental cue while the rest rely on cheaper social cues, imitating the movement of invested neighbors (Guttal and Couzin [2010], Fagan et al. [2012]). The collective migration problem can be abstracted to a collective tracking or learning problem, thus motivating the connection to nonlinear control design and in particular design that seeks to use minimal resources.

Social interactions in animal groups likewise play a central role in collective decision-making. The striking capabilities of animal groups as decision-makers also motivate the connection to nonlinear control design. Couzin et al. [2005] have shown that a group of animals moving together can make a critical collective decision on direction of motion, even if there is a conflict between directional preferences of informed individuals and uninformed individuals have no preference at all. No explicit signaling nor identification of informed individuals is assumed; individuals merely adjust their steering in response to socially acquired estimates on relative motion of neighbors. Another example is the collective decision-making of house-hunting honeybee swarms. Seeley and Buhrman [2001] showed how a unanimous best-of- N nest-site selection results from the dynamics of the “waggle dance,” an activity that honeybee scouts perform on the surface of the swarm to advertise the quality of potential nest sites. More recently, a cross-inhibitory signal among scouts that tempers the dancing has been shown to enable breaking deadlock among near-equal quality nest sites (Seeley et al. [2012]).

The present paper continues in Section 2 with some background and a very limited survey of relevant previous and current work. Models, mechanisms, bifurcation analyses, and prospects for nonlinear control design of multi-agent

systems are presented in Sections 3 and 4. Final remarks are provided in Section 5.

Section 3 examines the evolution of leadership in collective migration (tracking) using an adaptive dynamic model, motivated by the work of Guttal and Couzin [2010] and Torney et al. [2010]. Each agent has a scalar real-valued investment strategy that slowly adapts and changes the interaction network upon which the collective migration dynamics depend. Analytic results are presented in the case of very large populations using a mean-field approximation and in the case of relatively small populations using a time-scale separation. This section highlights the PhD research of Darren Pais (Pais [2012]).

Section 4 examines three models of collective decision-making. The first model, motivated by Couzin et al. [2005] and described in Section 4.1, uses a continuous-time model to derive rigorous conditions for decision versus poor compromise in the case of subsets of individuals in the group with conflicting preferences. This section builds on the PhD research of Benjamin Nabet (Nabet [2009]). The second model, motivated by the honeybee nest site selection dynamics of (Seeley et al. [2012]) and described in Section 4.2, is used to prove bifurcations in decision-making behavior, including deadlock breaking, as a function of the strength of the cross-inhibitory signal and the quality of the nest sites. The third model, described in Section 4.3 uses the replicator-mutator equations (Nowak et al. [2001]) to represent imitation and error in decision-making among finite alternatives. Conditions are proved on the network structure and error rate that yield stable limit cycles, which represent cycling among a subset of alternatives. The dynamics of the second and third models describe the distribution of very large populations and evolve on simplices. These two subsections highlight the PhD research of Darren Pais (Pais [2012]).

2. BACKGROUND

There is a wide range of high performing collective dynamics of animal groups, and we do not try to survey them all here. Indeed, collective dynamics can vary dramatically across species since different aggregations form as a result of very different selective pressures on self-interested individuals (Parrish and Hamner [1997]). In the field of collective animal behavior the mechanistic approach, which examines how group level behavior is produced, is tightly coupled to the functional approach, which seeks to explain why group level behaviors evolved in terms of natural selection (Sumpter [2010]). Couzin and Krause [2003] reviewed early progress on using mathematical modeling and self-organization theory to identify underlying principles of collective animal behavior. Vincent and Vincent [2000] examined the connections between the evolutionary process and stability and optimization in control design. Wei et al. [2009] evaluated feedback controlled steering strategies for pursuit in planar pursuit-evasion dynamics using an evolutionary game theoretic setting.

Social interactions in groups such as birds and fish are often modeled as repulsion, attraction, and/or alignment movements by individuals in response to nearby conspecifics that are, respectively, a little too close, a little too far, or heading off in a somewhat different direc-

tion. The numerical investigation of Couzin et al. [2002] showed how slow changes in the social interaction model parameters lead to transitions between parallel motion, circular motion and disordered swarming of the group. They showed further that for variation in the range of neighbors used in the alignment response, there is hysteresis in the transitions. Conditions for bi-stability and hysteresis of these motions were proved for a continuous-time dynamical model in Paley et al. [2007b].

It has been shown that animal groups perform well at tasks such as foraging and search, even in the presence of disturbances and noisy measurements. For example, individual aquatic animals during foraging may be limited to sampling their environment at scales much smaller than the scales at which resources are distributed. Furthermore, small-scale stochastic fluctuations may corrupt local gradients of resource. Grünbaum [1998] showed that schooling behavior enhances the group's ability to navigate and climb gradients. When individuals respond to social cues by staying close to one another and aligning their direction of motion, the corrupting effect of noise in individual measurements is dampened.

Motivated by these kinds of insights, Bachmayer and Leonard [2002] designed bio-inspired decentralized feedback strategies for a small formation of mobile sensor-equipped vehicles to climb the gradient in a sampled field. This approach was further developed into a more general methodology for gradient climbing (Ögren et al. [2004]) and level-set tracking (Zhang and Leonard [2010]) with adaptive formation shape changes for optimized filtering. The methodology was successfully demonstrated on autonomous underwater vehicles in an ocean sampling field experiment in Monterey Bay, California in 2003 (Fiorelli et al. [2006]). Torney et al. [2009] showed further that when individual fish continuously adjust how much they use social cues according to the confidence they have in their own strategy, the school can successfully track an advected chemical filament. Wu and Zhang [2012] have capitalized on this mechanism to prove a highly efficient source-seeking control strategy.

The model of social interactions in flocks and schools motivated development of a methodology for systematic stabilization of a family of parallel and circular motion patterns (Sepulchre et al. [2007, 2008]). The dynamics, which exploit a spatial extension of coupled oscillator dynamics, provide decentralized control strategies. This methodology was used both for studying the dynamics of schooling fish from experimental data (Paley et al. [2007a]) and for controlling the dynamics of mobile sensor networks into patterns that maximize information in data collected (Leonard et al. [2007]). The coordinated motion pattern design methodology was used successfully on six autonomous underwater gliders for 24 days straight in an adaptive ocean sampling field experiment in Monterey Bay, California in 2006 (Leonard et al. [2010]). Motivated by experimental data of schooling in which the fish exhibited coordinated oscillations in their speed, Swain and Leonard [2009] expanded the family of stabilizable coordinated motion patterns to include richer circular patterns that yield high spatial density in measurements, periodic

exchange of roles among agents, and possibilities for leveraging redundancy of sensor platforms.

Collective animal behavior was an early inspiration for collective robotics (Beckers et al. [1994]) and influenced the behavioral-based approach of the 1990's (Mataric [1992]). In recent years there has been renewed interest in enabling a collective capability in robotic groups using mechanisms attributed to social insect swarms (Berman et al. [2011], Trianni et al. [2011], Pratt and Sasaki [2012]). Butail et al. [2013] described the reconstruction of 3D trajectories from data of wild mosquito swarms and its use in informing a mechanistic model of male coordinated behavior. Chicoli et al. [2013] described reconstruction of 3D interactions of schooling fish and its use in a rigorous examination of the role of the flow field on the information transmission dynamics among the interacting fish. This work has potential implications for coordinated control of multi-agent systems in the air or water where there may be a steady flow.

3. COLLECTIVE MIGRATION

In the collective migration problem, a group of animals must learn and travel along a route that takes it from its current location to its new habitat or feeding ground. Migration is possible because animals can detect directional cues from the environment; however, acquiring this information may be costly in terms of time and energy, which could be spent on other demands such as reproduction, growth, or predator vigilance (Dall et al. [2005], Guttal and Couzin [2010]). Animals that migrate in groups can also use social cues such as observations of the relative movement of nearby animals, which may require little investment in time and energy. Collective migration will thus depend on there being a set of individuals in the group that invests in the costly environmental cue and can "lead" the rest of the group successfully.

Guttal and Couzin [2010] used simulation to study the evolutionary dynamics of a socially interacting migratory population where individuals adjust how much they invest in the costly environmental cue and how much they leverage social cues to optimize the trade-off between their migration benefit and their investment cost. They showed that for a sufficiently high cost, the population splits (branches) into invested "leaders" and uninvested "followers". They also predicted a strong hysteretic effect in which migration is lost at very high cost after which recovery of migration requires a significant reduction in cost.

The collective migration problem can be abstracted as a decentralized collective tracking (or learning) problem in which a multi-agent system as a whole should track (learn) an unknown signal. It is assumed that each agent has the capacity to measure the relative state of other agents (its neighbors), and to measure the unknown signal at a cost. In the case of an engineered system such as a network of mobile robotic sensors, measuring the unknown signal might involve turning on sensors that are power intensive such as a camera, as compared to less costly operation of proximity sensors for measuring the relative position of nearby robots. All measurements are assumed to be noisy, and thus performance can be defined both for an agent

and for the whole system in terms of the trade-off between tracking accuracy and investment cost. Other performance criteria might include tracking speed and adaptability to changes in the investment cost, the unknown signal, or the structure of the interconnections.

The problem of selecting a finite set of leaders that minimizes total system error in a multi-agent linear, stochastic, decentralized coordinated control dynamic has received recent attention in the literature (Patterson and Bamieh [2010], Clark et al. [2012], Lin et al. [2013]). In these works a fixed set of leaders is selected using a centralized (top-down) algorithm given a fixed interconnection graph. In contrast the evolution of collective migration describes dynamics in which the leadership and the interconnection graph emerges through a decentralized (bottom-up) and adaptive process. These adaptive dynamics are nonlinear even if the dynamics with fixed leadership are linear because the leadership investments and weights on social information are changing. In collaboration with Darren Pais, we analyzed the nonlinear dynamics and bifurcations to provide a means to systematically predict the influence of model parameters on the emergence of leadership and on collective tracking performance. This in turn suggests systematic approaches to high performance and efficient design of multi-agent system dynamics.

In Pais and Leonard [2013] the multi-agent system dynamics were modeled to approximate the detailed evolutionary simulation model used by Guttal and Couzin [2010]. The model is presented here using the terminology of collective migration; the abstraction to the collective tracking or learning problem is implied. There are N agents and the migration route is defined by a scalar direction μ . The state of agent i is defined by its migration direction¹ $x_i \in \mathbb{R}$ and its investment gain $k_i \in [0, 1]$. The dynamics have two time-scales: the migration dynamics of x_i are fast and the evolutionary dynamics of k_i are slow.

Our model of migration dynamics builds on the model used by Torney et al. [2010] in which individuals respond to the noisy measurement of the environmental cue with a gain that depends on k_i and to the noisy measurement of the social cues with a gain that depends on $(1 - k_i)$. The response to the environmental cue is modeled as an Ornstein-Uhlenbeck process (Uhlenbeck and Ornstein [1930]). The response to social cues is modeled as noisy linear consensus dynamics using a directed graph (digraph) representation of the interaction topology (Xiao et al. [2007]); this generalizes the model of Torney et al. [2010] which uses a mean-field approximation to reduce the social graph to an all-to-all interconnection.

For each agent i we denote the standard deviation of the additive white noise associated with the environmental measurement as σ_D and with the social cues as σ_{Si} . Consider a digraph \mathcal{G} that encodes the available social interconnections; we refer to the set of agents that agent i can sense as the neighbors of agent i . Let L be the Laplacian matrix associated with the social digraph. For convenience we define the change of coordinates $\tilde{x}_i = (x_i -$

¹ For moderate deviations of x_i from μ (e.g., due east), we can let $x_i \in \mathbb{R}$. This is not the case in Section 4.1 where there are two desirable directions (e.g., due east and due west) and directional dynamics on the circle are relevant.

$\mu)/\sigma_D$. Then, for each agent i the fast time-scale migration dynamics are

$$d\tilde{x}_i = -k_i^2 \tilde{x}_i dt - (1 - k_i)^2 L_i \tilde{\mathbf{x}} dt + \sqrt{k_i^2 + (1 - k_i)^2 \frac{\sigma_{Si}^2}{\sigma_D^2}} dW_i, \quad (1)$$

where L_i is the i th row of L and dW_i is the standard Wiener increment. The first term on the right of (1) drives x_i to μ at a rate k_i^2 . The second term on the right drives x_i to a weighted average of the states of its neighbors at a rate proportional to $(1 - k_i)^2$.

The third term of (1) is the sum of the two sources of noise. Inspired by the model of Torney et al. [2010] we let

$$\frac{\sigma_{Si}^2}{\sigma_D^2} = \beta^2 (1 - k_{nbhd,i}), \quad (2)$$

where $\beta > 0$ is a constant and $k_{nbhd,i}$ is the average investment gain of the neighbors of agent i .

In the fast migration time-scale, we can take k_i , $i = 1, \dots, N$ to be constant and the migration model (1)-(2) is then a linear stochastic differential equation. Rewriting as a matrix equation we get

$$d\tilde{\mathbf{x}} = -(K_1 + K_2 L) \tilde{\mathbf{x}} dt + S d\mathbf{W}, \quad (3)$$

where $K_1 = \text{diag}(k_i^2)$, $K_2 = \text{diag}((1 - k_i)^2)$ and $S = \text{diag}(\sqrt{k_i^2 + \beta^2 (1 - k_i)^2 (1 - k_{nbhd,i})})$.

In Pais and Leonard [2013] it is proved that for digraph \mathcal{G} , the deterministic dynamics of (3) are stable if and only if $k_j > 0$ for each j in a minimal root set² $\mathcal{R}(G)$. These conditions imply for the stochastic dynamics (3) that $\lim_{t \rightarrow \infty} E[\tilde{\mathbf{x}}(t)] = \mathbf{0}$, i.e., the steady-state expected value of x_i is μ for all i . In other words, if the set of leaders (invested agents) contains a minimal root set $\mathcal{R}(G)$, then the group will migrate. Further, the steady-state covariance matrix $\Sigma = \lim_{t \rightarrow \infty} E[\tilde{\mathbf{x}}(t)^T \tilde{\mathbf{x}}(t)]$ will satisfy the Lyapunov equation:

$$(K_1 + K_2 L) \Sigma + \Sigma (K_1 + K_2 L)^T = S S^T. \quad (4)$$

The migration accuracy of agent i is defined by the diagonal element $\sigma_{ss,i}^2 = \Sigma_{i,i}$, which is the steady state variance of x_i about μ . Variance about consensus was used in Young et al. [2013] to evaluate robustness to uncertainty in the collective behavior of starling flocks from data taken in the wild.

The fitness of agent i is defined following Guttal and Couzin [2010], Torney et al. [2010] as

$$F_i = \exp\left(\frac{-\sigma_{ss,i}^2}{2}\right) \exp(-ck_i^2), \quad (5)$$

where the first exponential models the migration speed of agent i and thus the benefit to agent i , and the second exponential models the cost to agent i associated with its investment in tracking. Here, $c > 0$ is a scaling cost parameter. This model of fitness, which can also be interpreted as a utility or payoff function, has features that are more generally representative of cooperative games. Do et al. [2010] argue for a general form that is the difference

² We define a minimal root set $\mathcal{R}(G)$ of a digraph \mathcal{G} to be a set of nodes of minimal cardinality such that there is a directed path from every node in \mathcal{G} to at least one node in $\mathcal{R}(G)$.

between a benefit that is sigmoidal (saturating) in the strategic parameter and a cost that is super linear in the strategic parameter. The fitness (5) has this general form since for $\beta > 1$ the migration speed is sigmoidal in k_i and the cost is quadratic in k_i .

The slow time-scale evolutionary dynamics of the investment strategies k_i can be computed using the roulette-wheel selection algorithm of Mitchell [1998] plus a small mutation on each “generation” of a large population: agents “reproduce” with a probability proportional to their fitness relative to the average fitness, and their “offspring” inherit their traits modulo a small mutation with a zero-mean Gaussian distribution. In the engineering context, this can be translated as a dynamic on the distribution across the group of control strategies over time driven by performance associated with these strategies (encoded by fitnesses).

An agent without any social interactions will evolve its strategy k_i to balance the increase in migration speed with increasing k_i against the decrease in investment cost with decreasing k_i ; for such an agent in isolation the strategy that maximizes its fitness is $k_i = (\sigma_D^2/8c)^{1/3}$. In the social context, however, the fitness landscape is frequency and network dependent, i.e., each agent’s fitness depends on how frequent within the group and where in the network are the different values of the different agents’ fitnesses; this implies an evolutionary game. For example, agents with strongly invested neighbors can develop high fitness without having to invest. As shown in the simulation studies of Guttal and Couzin [2010] the investment cost parameter c is critical to the evolutionary outcome. We have made this rigorous for our migration plus evolutionary dynamic model by proving bifurcations in the behavior of the group with c the bifurcation parameter (Pais and Leonard [2013]).

Let us first consider the case of a large population in which every agent can in principle sense every other agent³, i.e., the underlying graph is all-to-all with weights on edges dependent on the changing values of strategies k_i . The evolved behaviors of the two time-scale system are represented by the equilibrium solutions k_i^* of the slow evolutionary dynamics of the strategies k_i . A monomorphic solution corresponds to $k_i = k^*$ for all $i = 1, \dots, N$. Bistability of equilibrium values $k_f^* \approx 0$ and $k_l^* \approx 1$ represents the possibility of co-existence of two subpopulations, one of leaders with $k_i = k_l^*$ and one of followers with $k_i = k_f^*$.

In Pais and Leonard [2013] the existence and stability of the equilibrium strategies as a function of c was computed using an analytic computation of fitness from (4)-(5) and the analysis tools of evolutionary adaptive dynamics (Geritz et al. [1997a,b], Diekmann [2004]). Consider that most of the population, called the resident population, has common strategy k_R and let there be a very small mutant population with common strategy k_M . The relative fitness S of the small mutant population with respect to the larger resident population can be computed explicitly as

$$S(k_R, k_M) = F_M(k_R, k_M) - F_R(k_R),$$

where resident fitness F_R and mutant fitness F_M are computed from (4)-(5) as

$$F_R = \exp\left(-\frac{k_R^2 + \beta^2(1 - k_R)^3}{4(2k_R^2 - 2k_R + 1)} - ck_R^2\right),$$

$$F_M = \exp\left(-\frac{k_M^2 + \beta^2(1 - k_R)(1 - k_M)^2}{4(2k_M^2 - 2k_M + 1)} - ck_M^2\right).$$

Following Geritz et al. [1997a,b], Diekmann [2004], the evolutionary dynamics of the resident strategy k_R are

$$\frac{dk_R}{dt} = \gamma \left. \frac{\partial S}{\partial k_M} \right|_{k_M=k_R} =: \gamma g(k_R),$$

where $\gamma > 0$. An equilibrium k^* such that $g(k^*) = 0$ satisfies

$$k^* [\beta^2(1 - k^*) - 1] (k^* - 1) + 4ck^*(2k^{*2} - 2k^* + 1)^2 = 0.$$

A solution k^* is a convergent stable strategy (CSS) if $\left. \frac{dg}{dk_R} \right|_{k_R=k^*} < 0$, and either a local evolutionary stable strategy (ESS) or if $\left. \frac{\partial^2 S}{\partial k_M^2} \right|_{k_M=k_R=k^*} > 0$ a branching point.

Suppose that $\beta > 1$. Then, the follower solution $k_f^* = 0$ is an equilibrium for all values of c : it is unstable for $0 < c < c_0 = (\beta^2 - 1)/4$ and stable for $c > c_0$. In the range $0 < c < c_0$ there is only one other solution $k_l^*(c) > 0$, which decreases with increasing c , and it is ESS. Thus for relatively low cost $0 < c < c_0$, all agents will invest equally in the environmental signal, $k_i = k_l^*$. In the range $c_0 < c < c_1$ there is bi-stability of the leader solution k_l^* and the follower solution k_f^* with a third unstable solution; the bi-stability allows for the coexistence of leaders and followers. This is also true for $c_1 < c < c_2$; however, in this range k_l^* is a branching point which leads to an initially monomorphic population splitting into leaders and followers. For $c > c_2$ the only solution is k_f^* and thus migration is lost. After migration is lost at $c > c_2$, migration can only be recovered if c is reduced below c_0 when the follower solution becomes unstable again (see Pais and Leonard [2013] for expressions for c_1 and c_2 and further discussion). This bifurcation analysis describes, for large populations with dense interconnection, the changing behavior with increasing investment cost, from a single investment strategy to leader-follower strategies to loss of leadership and thus loss of migration at very high cost with the associated hysteresis as observed in the original simulations of Guttal and Couzin [2010].

How the behavior changes with the density of edges and the topology of the social graph is also critical to understanding the mechanisms of leadership and collective migration. Analyzing the differential fitness landscape for large populations with limited underlying social graphs can be challenging; however, parametric numerical studies of the evolutionary dynamics can be pursued relatively easily because fitness is easily computed from (4)-(5) even for N large. In Pais and Leonard [2013] evolutionary branching was studied as a function of probability of a pairwise edge (number of neighbors) in random graphs and cyclic lattices; it was shown that branching occurs above a minimum threshold in edge probability (number of neighbors) and this minimum can be quite low. That is, it only takes a small amount of connectivity, when

³ If agent i is fully invested, i.e., $k_i = 1$, then it does not pay attention to any other agents even if the social cues are available. Thus, even if the underlying social structure can be represented by an undirected all-to-all graph, distribution in strategies k_i will lead to a directed graph.

investment cost is high, for a population to split into leaders that invest and followers that rely on social cues.

The influence of the graph topology on emergence of leadership and collective migration can be further studied for small populations by replacing the evolutionary dynamics with a slow time-scale greedy adaptive dynamic such as

$$\dot{k}_i = \frac{\partial F_i}{\partial k_i}. \quad (6)$$

This can be interpreted as a decentralized adaptive control law for agent i , also known as gradient play (Shamma and Arslan [2005]). Assuming that agent i can measure its own fitness F_i it does not need to know the strategies of its neighbors. Alternatively, if it knows the strategies of its neighbors it could estimate its fitness F_i . The nonlinear dynamics of the adaptive dynamic network can be studied by substituting the expression for F_i from (4)-(5) into (6).

In the case of $N = 2$ a complete bifurcation analysis was computed with bifurcation parameter c , analogous to the large population case. In this case for low values of c , there is one stable equilibrium $k^*(c)$ that decreases with increasing cost c , i.e., both agents invest the same amount. With further increase in c there is first two saddle node bifurcations in which a symmetric pair of stable solutions appear; these correspond to one of the agents investing a lot and one investing very little. At a slightly higher c , there is a pitchfork bifurcation and the common solution k^* becomes unstable and thus the only stable solutions correspond to one leader and one follower. In similar analysis with larger N , there is multi-stability of solutions in which a subset of agents lead and the rest follow; as c increases the number of leaders decreases. In the case of an underlying star graph, it emerges that at low cost the fringe agents lead, at intermediate cost all agents invest but at a lower level, and at high cost only the center agent invests and the fringe agents follow. These and other results are described in Pais and Leonard [2013].

These bifurcation analyses can inform development of decentralized control strategies for multi-agent systems. For example, for adaptation of strategies k_i in collective tracking that derives as above, parameters such as c could be designed to target a desired number of leaders. An important step is deriving the relationship between the strategic outcomes and the group-level performance (e.g., using results on top-down optimal leader selection). Then, for example, fitness functions could be shaped so that the emergent locations of the leadership in the network yield optimal group performance. Alternatively, the fitness functions could be designed so that agents add or subtract communication links and the network changes structure to best leverage a prescribed set of leaders, for example, vehicles that carry extra sensors (such as cameras) to measure external signals. Approaches from distributed optimization (e.g., Nedic and Ozdaglar [2009]) may prove useful in this context.

4. COLLECTIVE DECISION-MAKING

In the collective decision-making problem a group of animals must make a choice among a finite number of alternatives, e.g., a set of possible resource-rich locations or a set of candidate nest sites. A unanimous decision

will typically be advantageous since individuals that get separated from the group will lose the benefits of group living. Sometimes there will be a best alternative, but a lesser alternative may often be better than a deadlock or a poor compromise. This collective decision-making problem translates to engineered multi-agent systems that must choose as a group among uncertain alternatives in tasks such as search and allocation.

In this section three different models and mechanisms for collective decision-making are presented with an emphasis on understanding behavior through bifurcation. The first model with agent-based dynamics similar to Section 3 is used to explore the role of neighborhood size and informed versus uninformed populations on decision-making between two alternative desirable directions. The second and third models both assume very large populations and formulate the dynamics of the distribution of the population among alternative choices. The second model is used to examine the role of recruitment in making the best collective decision and the role of cross-inhibition in breaking deadlock. The third model is an evolutionary model that includes replication and mutation and allows for the possibility of limit cycle behavior. This model, which resembles models of imitation and error, is used to explore the role of network topology and mutation (error) rate on stability and existence of equilibria and limit cycles.

4.1 Leadership

(Couzin et al. [2005]) used simulation to study how a group of agents moving together in the plane can make a collective decision for one of two equal value alternative directions of motion given conflict in the directional preferences of two small subgroups of informed agents. Inter-agent communication was limited to agents sensing the relative position and heading of nearby agents without being able to distinguish between the informed and uninformed. It was shown that a collective decision to move in one of the two alternative directions is made with high probability for sufficiently large magnitude preference conflict, i.e., large enough difference in preferred directions. Otherwise, the group makes a poor compromise and moves together in the direction that splits the difference between the two alternatives.

In collaboration with Tian Shen, Benjamin Nabet, Luca Scardovi, Iain Couzin and Simon Levin (Leonard et al. [2012]), we used a system of ordinary differential equations to approximate the high degree-of-freedom model of Couzin et al. [2005] and to rigorously explain the collective decision-making result and explore the parameter space. We present the model and summarize the main results here. There are N agents and each agent is a member of one of three subgroups: there are N_1 informed agents in subgroup 1 with a preference for moving in the direction θ_1 , there are N_2 informed agents in subgroup 2 with a preference for moving in the direction θ_2 , and there are N_3 uninformed agents in subgroup 3 with no preference on direction. Each agent j moves like a self-propelled particle in the plane at constant speed v_c in the direction $\theta_j(t)$ at time t such that its velocity in the plane is $\mathbf{v}_j(t) = (v_c \cos \theta_j(t), v_c \sin \theta_j(t))$. Each agent j puts a

weight $0 \leq a_{jl}(t) \leq 1$ at time t on its response to its measurement of the relative direction of motion of agent l ; if agent j cannot sense agent l then $a_{jl} = 0$. The neighbors of agent j at time t are the agents l such that $a_{jl}(t) > 0$.

The steering rate for each agent in subgroup 1 is

$$\frac{d\theta_j}{dt} = \sin(\bar{\theta}_1 - \theta_j(t)) + \frac{K_1}{N} \sum_{l=1}^N a_{jl}(t) \sin(\theta_l(t) - \theta_j(t)), \quad (7)$$

in subgroup 2 it is

$$\frac{d\theta_j}{dt} = \sin(\bar{\theta}_2 - \theta_j(t)) + \frac{K_1}{N} \sum_{l=1}^N a_{jl}(t) \sin(\theta_l(t) - \theta_j(t)), \quad (8)$$

and in subgroup 3 it is

$$\frac{d\theta_j}{dt} = \frac{K_1}{N} \sum_{l=1}^N a_{jl}(t) \sin(\theta_l(t) - \theta_j(t)). \quad (9)$$

The first term on the right of (7) and (8) steers the agent towards its preferred direction and the second term and only term of (9) steers the agent towards its neighbors' direction. The constant parameter $K_1 > 0$ weights the attention paid to social cues relative to preference. The dynamics of the interconnection weights are

$$\frac{da_{jl}}{dt} = K_2(1 - a_{jl}(t))a_{jl}(t)(\rho_{jl}(t) - r), \quad (10)$$

where $K_2 > 0$ is a constant, $\rho_{jl} = |\cos(\frac{1}{2}(\theta_j - \theta_l))|$ measures synchrony of direction of motion of l and j , and $0 \leq r \leq 1$ is a chosen fixed threshold representing an agent's sensing range. It holds that $\rho_{jl} = 1$ if l and j move in the same direction and $\rho_{jl} = 0$ if they move in opposite directions. If $\rho_{jl} > r$, then j and l are considered close enough to sense each other so a_{jl} eventually converges to the maximum interaction strength of 1. If $\rho_{jl} < r$, then j and l are not considered close enough to sense each other so a_{jl} eventually converges to 0.

The model exhibits fast and slow dynamics even for moderate values of K_1 and K_2 ; Leonard et al. [2012] proved a time-scale separation for large K_1 and K_2 . In the fast dynamics all agents within a subgroup quickly converge to a single direction and all a_{jl} converge quickly either connecting or disconnecting the subgroups. For example, $a_{jl} = A_{12}$ for all j in subgroup 1 and all l in subgroup 2, and subgroup 1 and subgroup 2 are connected if $A_{12} = 1$ and disconnected if $A_{12} = 0$. The slow variables are Ψ_1 , Ψ_2 , and Ψ_3 , where Ψ_k is the average direction of motion of subgroup k . The fast dynamics have a number of isolated solutions including eight that can be classified by A_{12} , A_{13} and A_{23} each taking value 0 or 1. The slow dynamics derived on each of these eight manifolds are

$$\begin{aligned} \frac{d\Psi_1}{dt} &= \sin(\bar{\theta}_1 - \Psi_1(t)) + \frac{K_1}{N} (A_{12}N_2 \sin(\Psi_2(t) - \Psi_1(t)) \\ &\quad + A_{13}N_3 \sin(\Psi_3(t) - \Psi_1(t))) \\ \frac{d\Psi_2}{dt} &= \sin(\bar{\theta}_2 - \Psi_2(t)) + \frac{K_1}{N} (A_{12}N_1 \sin(\Psi_1(t) - \Psi_2(t)) \\ &\quad + A_{23}N_3 \sin(\Psi_3(t) - \Psi_2(t))) \\ \frac{d\Psi_3}{dt} &= 0 + \frac{K_1}{N} (A_{13}N_1 \sin(\Psi_1(t) - \Psi_3(t)) \\ &\quad + A_{23}N_2 \sin(\Psi_2(t) - \Psi_3(t))). \end{aligned} \quad (11)$$

Three manifolds of interest are \mathcal{M}_{010} , which corresponds to $A_{12} = 0$, $A_{13} = 1$, $A_{23} = 0$ and a decision for direction $\bar{\theta}_1$, \mathcal{M}_{001} , which corresponds to $A_{12} = 0$, $A_{13} = 0$, $A_{23} = 1$ and a decision for direction $\bar{\theta}_2$, and \mathcal{M}_{111} , which corresponds to $A_{12} = 1$, $A_{13} = 1$, $A_{23} = 1$ and a solution between the directional choices, i.e., a poor compromise is made. It was shown that \mathcal{M}_{010} and \mathcal{M}_{001} , i.e., decisions for one of the two preferred directions, are both stable if and only if the difference in preferred direction is greater than the critical angle $\bar{\theta}_c = \cos^{-1}(2r^2 - 1)$; otherwise, the compromise solution on \mathcal{M}_{111} is stable.

This bifurcation in behavior as a function of the magnitude of the conflict qualitatively matches the original results of Couzin et al. [2005]. Further, it provides a sharp condition on the bifurcation point and the dependence of the bifurcation point on the sensing region parameter r . As r increases, the sensing region decreases and the critical value $\bar{\theta}_c$ decreases, meaning that it becomes easier for the group to make a decision. When r is small, the sensing region is large and the agents cannot distinguish between the alternatives; they function more as a filter and move according to a compromise between the two alternatives.

The model was also used to examine rigorously the role of the uninformed population size N_3 on the group decision-making dynamics. Leonard et al. [2012] showed how the model predicts an increasing likelihood of a collective decision for one or the other alternative with increasing N_3 . In a related study Couzin et al. [2011] examined the beneficial role of the uninformed population in preserving democratic consensus in the case of conflict where the smaller informed subgroup has a stronger preference than the larger informed subgroup.

The rigorously derived relationship between decision-making performance and parameters such as sensing range and relative size of informed and uninformed populations suggests opportunities for design of decentralized control strategies for high performance collective decision-making in engineered networks.

4.2 Recruitment and Cross-Inhibition

In the honeybee house-hunting problem, an entire swarm must unanimously choose a nest site where it will live as a new colony with its queen. The process starts with a number of bees each scouting out a possible nest site and ends with the swarm choosing the best of the scouted out alternatives. In contrast to the limitations on signaling in groups described in Sections 3 and 4.1, each informed honeybee scout uses explicit signaling in the form of a "waggle dance" to recruit uninformed bees to commit to its discovered nest site. Seeley et al. [2012] showed from empirical data in which there were two alternative sites that the scouts also use a cross-inhibitory stop-signal to make the scouts recruiting for the competing alternative stop dancing. This stop-signal contributes positively to the collective decision-making; of particular note, it was shown to help with breaking deadlock in the case of two near-equal value alternative sites.

Seeley et al. [2012] derived a model of the mean-field population-level dynamics under the assumption that the total bee population size N is very large. The model

describes the dynamics of $y_A = N_A/N$ and $y_B = N_B/N$, the changing fraction of bees committed to nest site A and B , respectively. Since N is constant, $y_A + y_B + y_U = 1$ where y_U is the fraction of uncommitted bees. Thus, the dynamics evolve on the two-dimensional simplex as

$$\begin{aligned}\dot{y}_A &= \gamma_A y_U - y_A(\alpha_A - \rho_A y_U + \sigma_B y_B) \\ \dot{y}_B &= \gamma_B y_U - y_B(\alpha_B - \rho_B y_U + \sigma_A y_A),\end{aligned}$$

where γ_i is the rate of scouting discovery and commitment, α_i is the rate of abandonment of commitment, ρ_i is the rate of recruitment, and σ_i is the rate of stop-signaling. It is assumed that all but the stop-signal rate depend on the value v_i (quality) of the nest site; in particular, $\gamma_i = \rho_i = v_i$ and $\alpha_i = 1/v_i$. The two stop-signal rates are assumed to be the same, i.e., $\sigma_i = \sigma$. When y_A or y_B crosses a quorum threshold ω , a collective decision is reached.

Together with Darren Pais, Patrick Hogan and James Marshall, we have studied these nonlinear dynamics and identified bifurcations in the collective decision-making behavior in terms of the stop-signal rate σ , the mean value of the alternatives $\bar{v} = (v_A + v_B)/2$ and the difference in value of the alternatives $\Delta v = v_A - v_B$ (Pais et al. [2013]). For example, it can be shown that for near-equal value sites and a given rate of stop-signaling there is a threshold such that deadlock is broken for good quality nest sites, i.e., when the mean value \bar{v} is above the threshold, and deadlock is maintained for poor quality nest sites, i.e., when the mean value \bar{v} is below the threshold. The deadlock in the case of poor quality sites could be useful in case a newer better alternative is discovered. The bifurcation analyses provide a systematic way to understand how the stop-signal rate can be used to tune the threshold for breaking deadlock, to tune the minimum difference in value between the alternatives required for reliable discrimination, and to tune the speed-accuracy trade-off in the case in which uncertainty in the rates is taken into account.

It is nontrivial to design dynamics for a decision-maker to reliably choose the best among finite alternatives when they are sufficiently different and to quickly choose any of the alternatives when they are similar and of sufficiently good quality. The bifurcation analyses described above point to opportunities for collective decision-making with tunable deadlock breaking in engineered multi-agent systems by designing cross-inhibitory signals into decentralized control laws with gains that adapt with the changing environment.

4.3 Imitation and Error

The replicator-mutator equations from evolutionary dynamics model replication of a finite set of n strategies in terms of their relative fitnesses and mutation in terms of a probability of spontaneous switches of strategies (Bürger [1998], Page and Nowak [2002]). These dynamics have been used to model a variety of systems in biology (Komarova [2004]), including autocatalytic reaction networks (Stadler and Schuster [1992]), and the evolution of language, where strategies represent different grammars (Nowak et al. [2001]). They have also been used recently to model the

dynamics of social networks (Olfati-Saber [2007], Hussein [2009]) and wireless multi-agent networks (Tembine et al. [2010], Wang and Hussein [2010]). Under certain conditions these dynamics represent models of decision-making dynamics in networked multi-agent systems where agents decide among n choices: replication captures imitation of successful strategies and mutation captures errors.

Under the assumption of very large populations, the dynamics describe the evolution of x_i , the fraction of the total population with strategy i , for $i = 1, \dots, n$. Since $\sum_{i=1}^n x_i = 1$, the dynamics evolve on an $(n-1)$ -dimensional simplex. The fitness of strategy i is defined as $f_i = \sum_{j=1}^n b_{ij} x_j$, where $b_{ij} > 0$ is the payoff to an agent with strategy i on interacting with an agent with strategy j and $b_{ii} = 1$ for all i . The payoff matrix B with elements b_{ij} can be interpreted as the adjacency matrix of a directed graph. Mutation probabilities q_{ij} , which define the probability of a spontaneous switch from strategy i to j , are typically defined in terms of a mutation strength parameter $\mu \in [0, 1]$ and can depend on the b_{ij} and thus also adopt a graph structure. The replicator-mutator dynamics are given by

$$\dot{x}_i = \sum_{j=1}^n x_j f_j(\mathbf{x}) q_{ji}(\mu) - x_i \sum_{k=1}^n f_k(\mathbf{x}) x_k. \quad (12)$$

In the literature the analysis of the replicator-mutator equations (12) have focused on equilibrium behaviors and fitness matrices with a lot of symmetry. However, it has been shown that the symmetric cases are structurally unstable and that breaking symmetry can lead to chaotic dynamics (Komarova and Levin [2010], Mitchener and Nowak [2004]). Motivated by these results and an interest in examining and designing dynamics that are richer than equilibrium behavior, we have studied symmetry breaking in the replicator-mutator equations (12). In collaboration with Darren Pais and Carlos Caicedo, we proved bifurcations as a function of μ for asymmetric (directed) graphs (Pais et al. [2012]). We showed rigorously how certain asymmetries make possible a departure from behaviors described only by equilibrium states; indeed we proved conditions for the existence of stable limit cycles through Hopf bifurcations.

In the context of the imitation and error decision-making paradigm, the parameter μ can be interpreted as a learning parameter that decreases in value as the agents learn and do a better job at avoiding error. For large μ close to 1, the stable solution is the mixed equilibrium, where $x_i = 1/n$ for all i . As μ decreases (i.e., learning improves), one or more limit cycles appear which correspond to a majority of the population making cycles through a set of alternatives. As μ decreases further towards zero, the limit cycles disappear and the stable solutions approach the pure single strategy solutions, which each correspond to a unanimous choice for a single alternative.

This interpretation suggests opportunities for designing adaptive dynamics for the parameter μ in an engineered decision-making network such that μ decreases as the system gains information, explores options and learns, and then makes an informed collective choice among alternatives.

5. FINAL REMARKS

Animals that travel in a group exhibit remarkable collective behavior: as individuals they are limited in sensing, computation and actuation, but as a group they excel at tracking and decision-making tasks in uncertain and dynamic environments. We have shown ways in which nonlinear dynamics and control theory can be used to rigorously investigate mechanisms of feedback and interaction in this context. Our goal is twofold: we aim to help explain how complex collective animal behavior can emerge from relatively simply individual behaviors, and we aim to lay the foundations for nonlinear control design of engineered multi-agent system dynamics that inherit the remarkable capabilities of animal groups.

To better understand the adaptive behavior of animal groups and the corresponding opportunities for design of adaptive dynamics in a changing environment, we have analyzed bifurcations in the collective dynamics of groups as a function of system and environmental parameters. For example, we have studied bifurcations in the evolution of leadership and network weights in collective tracking as a function of a parameter c that models the cost of leadership to individuals. In the case of animal groups, c models the difficulty in sensing the environmental cue, including the extent of fragmentation in the environment. In the case of designed groups, c can be used as a design parameter that incentivizes leadership as appropriate for changing operational or environmental circumstances.

In problems of collective decision-making we have studied bifurcations that determine how the group can avoid poor compromise between alternatives as a function of sensing range or can tune the conditions for maintaining or breaking deadlock as a function of a cross-inhibitory signal among subgroups with commitment to competing alternatives. We have also studied bifurcations in a decision-making model as a function of a learning parameter and shown conditions for stable limit cycles, which may contribute to useful information gathering during the collective decision-making process. For design, the sensing range, cross-inhibitory strength, or learning parameter can each be used as a design parameter to control collective decision-making behavior in response to changes in alternatives, availability of resources, etc.

We have used models and methods from evolutionary dynamics to examine the evolution of feedback strategies, and we have discussed how these tools can be translated into adaptive control laws for engineered systems. Future work is needed to connect evolutionarily stable outcomes with optimal and robust emergent design solutions. It is also of interest to investigate variations and generalizations of the models and mechanisms of collective behavior described here, to explore additional exemplary collective dynamic phenomena and to evaluate performance bounds. An ultimate goal is to develop provable and systematic control design methodologies that yield high-performing, adaptive multi-agent system dynamics and address the myriad challenges that arise in the growing number of real-world applications.

6. ACKNOWLEDGMENTS

The author thanks Darren Pais, Dan Koditschek and Vaibhav Srivastava for helpful comments during the preparation of this paper.

REFERENCES

- P. Antsaklis and J. Baillieul, editors. *IEEE Transactions on Automatic Control: Special Issue on Networked Control Systems*, volume 49:9. IEEE, 2004.
- P. Antsaklis and J. Baillieul, editors. *Proceedings of the IEEE: Special Issue on Technology of Networked Control Systems*, volume 95:1. IEEE, 2007.
- R. Bachmayer and N.E. Leonard. Vehicle networks for gradient descent in a sampled environment. In *Proceedings of the IEEE Conference on Decision and Control*, pages 112–117, 2002.
- 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. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237, 2008.
- R. Beckers, O.E. Holland, and J.L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Proc. A-Life*. MIT Press, 1994.
- S. Berman, Q. Lindsey, M. S. Sakar, V. Kumar, and S. C. Pratt. Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems. *Proceedings of the IEEE*, 99(9):1470–1481, 2011.
- F. Bullo, J. Cortés, and B. Piccoli, editors. *SIAM Journal on Control and Optimization: Special Issue on Control and Optimization in Cooperative Networks*, volume 48:1. SIAM, 2009.
- R. Bürger. Mathematical properties of mutation-selection models. *Genetica*, 102-103:279–298, 1998.
- S. Butail, N. Manoukis, M. Diallo, J. M. Ribeiro, and D. A. Paley. The dance of male anopheles gambiae in wild mating swarms. *J. Medical Entomology*, 2013.
- A. Chicoli, S. Butail, J. Lun, J. Bak-Coleman, S. Coombs, and D. A. Paley. The effects of flow on fish school structure and startle response probability. *Preprint*, 2013.
- A. Clark, L. Bushnell, and R. Poovendran. A supermodular optimization framework for leader selection under link noise in linear multi-agent systems. *arXiv:1208.0946 [cs.SY]*, 2012.
- I. D. Couzin and J. Krause. Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, 32:1–75, 2003.
- I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1): 1–11, 2002.
- I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- I. D. Couzin, C. C. Ioannou, G. Demirel, T. Gross, C. J. Torney, A. Hartnett, L. Conradt, S. A. Levin, and N. E. Leonard. Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(1578), 2011.

- S. R. X. Dall, L.-A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens. Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution*, 20(4):187–193, 2005.
- O. Diekmann. A beginner’s guide to adaptive dynamics. *Banach Center Publications*, 63:47–86, 2004.
- A. L. Do, L. Rudolf, and T. Gross. Patterns of cooperation: fairness and coordination in networks of interacting agents. *New Journal of Physics*, 12:063023, 2010.
- W. F. Fagan, R. S. Cantrell, C. Cosner, T. Mueller, and A. E. Noble. Leadership, social learning, and the maintenance (or collapse) of migratory populations. *Theoretical Ecology*, 5:253–264, 2012.
- E. Fiorelli, N.E. Leonard, P. Bhatta, D. Paley, R. Bachmayer, and D.M. Fratantoni. Multi-AUV control and adaptive sampling in Monterey Bay. *IEEE J. Oceanic Engineering*, 31(4):935–948, 2006.
- J. M. Fryxell and A. R. E. Sinclair. Causes and consequences of migration by large herbivores. *Trends in Ecology & Evolution*, 3(9):237–241, 1988.
- S. A. H. Geritz, E. Kisdi, G. Meszéna, and J. A. J. Metz. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology*, 12(1):35–57, 1997a.
- S. A. H. Geritz, J. A. J. Metz, É. Kisdi, and G. Meszéna. Dynamics of adaptation and evolutionary branching. *Physical Review Letters*, 78(10):2024–2027, 1997b.
- D. Grünbaum. Schooling as a strategy for taxis in a noisy environment. *Evolutionary Ecology*, 12:503–522, 1998.
- V. Guttal and I. D. Couzin. Social interactions, information use, and the evolution of collective migration. *Proceedings of the National Academy of Sciences*, 107(37):16172–16177, 2010.
- R. A. Holland, M. Wikelski, and D. S. Wilcove. How and why do insects migrate? *Science*, 313(5788):794–796, 2006.
- I. Hussein. An individual-based evolutionary dynamics model for networked social behaviors. In *Proceedings of the American Control Conference*, pages 5789–5796, 2009.
- N. L. Komarova. Replicator-mutator equation, universality property and population dynamics of learning. *Journal of Theoretical Biology*, 230(2):227–239, 2004.
- N. L. Komarova and S. A. Levin. Eavesdropping and language dynamics. *Journal of Theoretical Biology*, 264(1):104–118, 2010.
- J. Krause and G. D. Ruxton. *Living in Groups*. Oxford University Press, 2002.
- N. E. Leonard, D. A. Paley, R. E. Davis, D. M. Fratantoni, F. Lekien, and F. Zhang. Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay. *J. Field Robotics*, 27(6):718–740, 2010.
- N. E. Leonard, T. Shen, B. Nabet, L. Scardovi, I. D. Couzin, and S. A. Levin. Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences*, 109(1):227–232, 2012.
- N.E. Leonard, D.A. Paley, F. Lekien, R. Sepulchre, D.M. Fratantoni, and R.E. Davis. Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE*, 95:48–74, 2007.
- F. Lin, M. Fardad, and M. R. Jovanovic. Algorithms for leader selection in stochastically forced consensus networks. *arXiv:1302.0450 [math.OC]*, 2013.
- M. J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In *2nd Int. Conf. Simulation of Adaptive Behavior*, pages 432–441. MIT Press, 1992.
- M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods for Multiagent Networks*. Princeton University Press, 2010.
- M. Mitchell. *An Introduction to Genetic Algorithms*. Bradford, 1998.
- W. G. Mitchener and M. A. Nowak. Chaos and language. *Proceedings of the Royal Society B: Biological Sciences*, 271(1540):701–704, 2004.
- B. Nabet. *Dynamics and Control in Natural and Engineered Multi-Agent Systems*. PhD thesis, Princeton University, Princeton, NJ, 2009.
- A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.
- M. A. Nowak, N. L. Komarova, and P. Niyogi. Evolution of universal grammar. *Science*, 291(5501):114–118, 2001.
- P. Ögren, E. Fiorelli, and N.E. Leonard. Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment. *IEEE Trans. Automatic Control*, 49:1292–1302, 2004.
- R. Olfati-Saber. Evolutionary dynamics of behavior in social networks. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4051–4056, 2007.
- K. M. Page and M. A. Nowak. Unifying evolutionary dynamics. *Journal of Theoretical Biology*, 219(1):93–98, 2002.
- D. Pais. *Emergent Collective Behavior in Multi-Agent Systems: An Evolutionary Perspective*. PhD thesis, Princeton University, 2012.
- D. Pais and N. E. Leonard. Adaptive network dynamics and evolution of leadership in collective migration. *Physica D, arXiv:1303.2242 [nlin.AO]*, 2013.
- D. Pais, C. H. Caicedo-Nunez, and N. E. Leonard. Hopf bifurcations and limit cycles in evolutionary network dynamics. *SIAM Journal on Applied Dynamical Systems*, 11(4):1754–1884, 2012.
- D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall. A mechanism for value-sensitive decision-making. Preprint, 2013.
- D. A. Paley, N. E. Leonard, R. Sepulchre, D. Grünbaum, and J. K. Parrish. Oscillator models and collective motion: Spatial patterns in the dynamics of engineered and biological networks. *IEEE Control Systems Magazine*, 27(4):89–105, 2007a.
- D. A. Paley, N. E. Leonard, R. J. Sepulchre, and I. D. Couzin. Spatial models of bistability in biological collectives. In *Proceedings of the IEEE Conference on Decision and Control*, 2007b.
- J. K. Parrish and L. Edelstein-Keshet. Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, 284(5411):99–101, 1999.
- J. K. Parrish and W. M. Hamner, editors. *Animal Groups in Three Dimensions*. Cambridge University Press, 1997.
- S. Patterson and B. Bamieh. Leader selection for optimal network coherence. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2693–2697, 2010.

- S. C. Pratt and T. Sasaki. Groups have a larger cognitive capacity than individuals. *Current Biology*, 22(19):827–829, 2012.
- T. D. Seeley and S. C. Buhrman. Nest-site selection in honey bees: how well do swarms implement the “best-of- N ” decision rule? *Behavioral Ecology and Sociobiology*, 49:416–427, 2001.
- T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N.R. Franks, and J. A. R. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–111, 2012.
- R. Sepulchre, D. A. Paley, and N. E. Leonard. Stabilization of planar collective motion: All-to-all communication. *IEEE Trans. Automatic Control*, 52(5):811–824, 2007.
- R. Sepulchre, D. A. Paley, and N. E. Leonard. Stabilization of planar collective motion with limited communication. *IEEE Trans. Automatic Control*, 53(3):706–719, 2008.
- J. S. Shamma and G. Arslan. Dynamic fictitious play, dynamic gradient play, and distributed convergence to nash equilibria. *IEEE Transactions on Automatic Control*, 50(3):312–326, 2005.
- P. F. Stadler and P. Schuster. Mutation in autocatalytic reaction networks. *Journal of Mathematical Biology*, 30(6):597–632, 1992.
- D. J. T. Sumpter. *Collective animal behavior*. Princeton University Press, 2010.
- D. T. Swain and N. E. Leonard. On the trajectories and coordination of steered particles with time-periodic speed profiles. In *Proceedings of the American Control Conference*, pages 1286–1291, 2009.
- H. Tembine, E. Altman, R. El-Azouzi, and Y. Hayel. Evolutionary games in wireless networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 40(3):634–646, 2010.
- C. J. Torney, Z. Neufeld, and I. D. Couzin. Context-dependent interaction leads to emergent search behavior in social aggregates. *Proceedings of the National Academy of Sciences*, 106(52):22055–22060, 2009.
- C. J. Torney, S. A. Levin, and I. D. Couzin. Specialization and evolutionary branching within migratory populations. *Proceedings of the National Academy of Sciences*, 107(47):20394–20399, 2010.
- V. Trianni, E. Tuci, K. M. Passino, and J. A. R. Marshall. Swarm cognition: an interdisciplinary approach to the study of self-organising biological collectives. *Swarm Intelligence*, 5:3–18, 2011.
- G. E. Uhlenbeck and L. S. Ornstein. On the theory of the Brownian motion. *Physical Review*, 36(5):823, 1930.
- T. L. Vincent and T. L. S. Vincent. Evolution and control system design: The evolutionary game. *IEEE Control Systems Magazine*, 20(5):20–35, 2000.
- Y. Wang and I Hussein. Evolutionary bandwidth allocation and routing in large-scale wireless sensor networks. In *Proceedings of the American Control Conference*, pages 1850–1855, 2010.
- E. Wei, E. W. Justh, and P. S. Krishnaprasad. Pursuit and an evolutionary game. *Proceedings of the Royal Society A*, 465(2105):1539–1559, 2009.
- W. Wu and F. Zhang. Robust cooperative exploration with a switching strategy. *IEEE Trans. Robotics*, 28(4):828–839, 2012.
- L. Xiao, S. Boyd, and S. J. Kim. Distributed average consensus with least-mean-square deviation. *J. Parallel and Distributed Computing*, 67(1):33–46, 2007.
- G. F. Young, L. Scardovi, A. Cavagna, I. Giardina, and N. E. Leonard. Starling flock networks manage uncertainty in consensus at low cost. *PLoS Computational Biology*, 9(1):1–7, 2013.
- R. A. Zachary, S. S. Sastry, and V. Kumar, editors. *Proceedings of the IEEE: Special Issue on Swarming in Natural and Engineered Systems*, volume 99:9. IEEE, 2011.
- F. Zhang and N. E. Leonard. Cooperative filters and control for cooperative exploration. *IEEE Trans. Automatic Control*, 55(3):650–663, 2010.