

NETWORK SCIENCE

Optimal network topology for responsive collective behavior

David Mateo^{1*}, Nikolaj Horsevad¹, Vahid Hassani¹, Mohammadreza Chamanbaz^{1,2}, Roland Bouffanais¹

Animals, humans, and multi-robot systems operate in dynamic environments, where the ability to respond to changing circumstances is paramount. An effective collective response requires suitable information transfer among agents and thus critically depends on the interaction network. To investigate the influence of the network topology on collective response, we consider an archetypal model of distributed decision-making and study the capacity of the system to follow a driving signal for varying topologies and system sizes. Experiments with a swarm of robots reveal a nontrivial relationship between frequency of the driving signal and optimal network topology. The emergent collective response to slow-changing perturbations increases with the degree of the interaction network, but the opposite is true for the response to fast-changing ones. These results have far-reaching implications for the design and understanding of distributed systems: a dynamic rewiring of the interaction network is essential to effective collective operations at different time scales.

INTRODUCTION

A wide range of complex systems are characterized by relatively simple dynamical rules while still producing excessively complex emergent collective behaviors. Examples abound in the natural world [e.g., a flock of birds, a school of fish, a swarm of insects (1–9)], in social systems [e.g., social networks (10–12)], and in engineered multi-agent systems [e.g., self-organized networks of mobile sensors, multi-vehicle coordination, and swarm robotics systems (13–16)].

Historically, particular attention has been directed toward investigating varieties of collective behaviors obtained by testing a wide range of local agent-to-agent interaction rules (6, 9). Collective behaviors have also been investigated from the network-theoretic perspective (4, 8, 17–21). It is now clear that such rich collective behaviors are the outcome of a complex interplay between network topology—characteristic of the group-level organization—and the dynamical laws at the agent's level (4, 8, 20–22).

Many collective behaviors can be studied through the lens of distributed consensus problems, including collective motion in animal groups and multi-robot systems. Over the past decade, the number of studies on decentralized consensus and cooperation in networked multi-agent systems has experienced a spectacular growth, with concomitant developments in various fields of engineering and science (2, 3, 23–26). Consensus dynamics is the cornerstone of cooperative control strategies for vehicular formation (13, 16, 23), swarm robotics (14, 15), and synchronization of coupled oscillators (23, 27). Decentralized consensus is also at the core of collective opinion dynamics and complex contagion processes in social networks (10–12), as well as complex collective responses in biological swarms (3–8).

Previous studies focused on establishing the influence of the interaction network topology on (i) the capacity of the collective to reach consensus in the presence of noise, communication constraints, and time delays (21, 23); (ii) the speed of consensus (i.e., its convergence rate) (18, 25, 28); (iii) the stability and stabilization of consensus (23); and (iv) the ability to steer the system toward a particular consensus value by means of various control techniques such as pinning control, cooperative tracking control, or model reference consensus (19, 20).

However, the effects of the network topology on other dynamical properties of distributed multi-agent systems such as their adaptivity or responsiveness to external perturbations have received considerably less attention (4).

It is important to emphasize that a capacity for fast consensus is not necessarily indicative of a responsive collective behavior. For instance, ferromagnets at low temperature exhibit a global spontaneous magnetization—a process that can be described by a distributed consensus protocol. It is known that both the degree of consensus (i.e., magnetization) and the speed at which it is reached increase with decreasing temperature, but the capacity of the system to respond to external perturbations is maximized at a finite critical temperature.

Similarly, in the context of animal collective motion, it has been observed that midges exhibit low levels of ordering while maintaining large connected correlations, thus having a high collective response (5). With these observations, the authors eloquently argued that one must be careful in relating collective order (i.e., degree of consensus) with the collective responsiveness. The collective response of the animal group was obtained experimentally by measuring the correlations in the fluctuations of their behavior. While inferring a collective response to external perturbations from these fluctuations is not formally justified for out-of-equilibrium systems, extensive numerical studies (29) have shown that this equivalence holds in the context of collective motion based on distributed heading consensus. Moreover, simulations have shown that this measure of susceptibility is a good indicator of the group's performance in biologically relevant functions such as predator avoidance (8). These facts along with other empirical evidence have led to the conclusion that responsiveness, rather than high consensus or order, is the true hallmark of collective behavior (3).

To study how the responsiveness of a collective is affected by its interaction network topology, we consider an elementary example of distributed decision-making: a linear time-invariant (LTI) system of agents performing consensus over a scalar state variable. The agents—i.e., the nodes of the interaction network—are all identical, except for one “leader” [also known as “stubborn” agent in some contexts (12, 25)] with some arbitrary predefined dynamics. From the control-theoretic perspective, this leader introduces a time-varying control input signal into the system. In the biological context, this dynamical leader represents a member of a swarm with access to privileged information about a food

Copyright © 2019
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

¹Singapore University of Technology and Design, 8 Somapah Road, Singapore 487372, Singapore. ²Arak University of Technology, Daneshgah Road, Arak, Iran. *Corresponding author. Email: david.mateo.valderrama@gmail.com

source or a threat, i.e., the dynamics of the leader can be considered to be a reaction to some external perturbation within the environment. For instance, in a school of fish, a single fish detecting the approach of a predator swiftly changes its direction of travel, and this rapid signal—the local external perturbation—propagates through the school, thereby triggering a large-scale evasive maneuver whose effectiveness is critical to the survival of the group.

Investigating the propagation of a local perturbation with a possibly broad spectrum of time scales is of critical importance to a vast breadth of decentralized networked systems, e.g., the fast shutdown at one end of a power grid can cascade into a large-scale blackout, a snowstorm at one critical node of an airport network generates delays throughout the entire system, and a fad introduced by an “influencer” propagates and amplifies through a scale-free social network. For artificial multi-agent systems, introducing a leader can facilitate a range of formation control techniques (13, 16) by means of pinning control or cooperative tracking control (19, 20, 26). In such a scenario, having a responsive collective behavior is crucial in the case that the target formation changes with time.

Here, the leader provides a gateway to injecting local perturbations in the emergent collective behavior. This allows us to study the ensuing collective response as a function of the time scale of the injected perturbation and the topology of the interaction network. A similar LTI framework was considered in (4) to study the effect of the network topology on the consensus. However, the study was limited to the long-time consensus dynamics arising from a global perturbation—in the form of white noise injected into the dynamics of all agents.

To quantify the collective responsiveness, several metrics are used across different communities—susceptibility, gain, or frequency response—but the definition is consistent: The response is measured by the rate of change in the variable of interest (in our case, the state variable participating in the consensus dynamics) with respect to a change in the external perturbation applied (in our case, the input signal injected through the leader). For a linear model, this rate of change can be expressed analytically as Eq. 5 and is a complex N -vector in the state space of agents (see Materials and Methods). We can quantify the collective frequency response of the system using the square of the norm of this vector H^2 . Because each component $h_i(\omega)$ of $\mathbf{H}(\omega)$ corresponds to the individual frequency response of a participating agent i , and given that $|h_i| \leq 1$, the collective frequency response $H^2(\omega)$ can be interpreted as the number of agents that are able to respond or follow the leader, when the dynamics of the latter varies with frequency ω . From a statistical perspective, H^2 measures the size of the fluctuations in the consensus state variable induced by a localized perturbation. For a connected topology, one expects $H^2 \approx N$ at low frequency, i.e., for frequencies below the system’s speed of consensus as determined by the smallest eigenvalue of the grounded Laplacian (25). In the limit of zero frequency, the system is at steady state and the consensus is guaranteed; therefore, $H^2(\omega = 0) = N$. At higher frequency, the collective response always decays with increasing ω , but the way that it decays is intricately dependent on the details of the connectivity between agents.

We show that the degree distribution of the interaction network is a central element in controlling the responsiveness of the collective. Even a relatively unadorned model, such as the linear consensus protocol, displays a rich phenomenology that crucially depends on the time scale of the changes in the input signal. Specifically, when the system is driven at low frequency, an increase in the number of interagent connections always improves the collective response. At high frequency, on the con-

trary, reducing the agents’ degree benefits the responsiveness. More generally, in the presence of a dynamic driving signal changing at a given time scale, it is possible to tune the interaction network to maximize the collective response. Last, we present experimental validations of our findings with a swarm of 11 robots performing heading consensus. We measure the capacity of the robots to align their orientation to a leader agent as a function of the degree distribution. From low to high frequency, the variations of the measured collective response with the degree of the interaction network are in very good agreement with the trend predicted by the distributed linear consensus theory.

These findings have far-reaching implications for the design of artificial swarms or interaction networks. They teach us that, to maximize the collective response, the system requires the ability to dynamically adjust the degree of the interaction network depending on the time scale at which it should respond.

RESULTS

Influence of the number of connections on the frequency response

In (8), it was shown that a breadth of networked systems can maximize their performance by tuning the number of connections to a specific finite value: A self-propelled particle swarm can increase its capacity to avoid a predator, and a collective performing distributed decision-making can react to an external influence faster. In the case of distributed consensus, it was seen that the collective frequency response of the system to a single leader when the agents are connected by a regular one-dimensional (1D) grid (a ring) is maximized for a number of neighbors k^* that depends on the frequency of the leader.

Figure 1 shows the collective frequency response of a system performing distributed linear consensus (Eq. 3) when one leader and $N = 2048$ agents are connected to their k -closest neighbors in a ring topology. When all agents are connected to each other ($k = N$), the response of the system is a first-order low-pass filter with a cutoff frequency of $1/N$. This all-to-all connectivity provides an optimal response for any frequency below $\omega_{\text{low}} \approx 2 \times 10^{-3}$. Above this threshold, a lower network degree yields a higher response.

The higher the frequency, the better low degrees (small k) respond as compared to high (large k) ones. For instance, $k = 30$ yields a higher response than $k = N$ for $\omega \geq 2.24 \times 10^{-3}$. In turn, the response of the system with $k = 10$ exceeds that of $k = 30$ for frequencies $\omega \geq 1.38 \times 10^{-2}$. A minimal, closest-neighbor connectivity ($k = 2$) provides the optimal response for any frequency above $\omega_{\text{high}} \approx 2.78 \times 10^{-1}$.

The optimal degree k^* corresponding to the highest response at a given frequency ω is displayed in Fig. 1 for several system sizes N . For large enough systems, the optimal degree follows a scaling law of the form

$$k^*(\omega) = K_0 \omega^{-\gamma} \quad (1)$$

with $K_0 = 1.56$ and $\gamma = 0.56$. The optimal degree does not scale with the size of the system N in any observable way. The finite-size effects manifest at low frequency, where k^* “jumps” from its bulk value given by Eq. 1 to $k^* = N$.

The frequency at which this jump occurs does depend on the size of the system and corresponds approximately to $k^* = \sqrt{N}$. In the limit $N \rightarrow \infty$, a finite connectivity is always preferable over the all-to-all topology for any finite frequency.

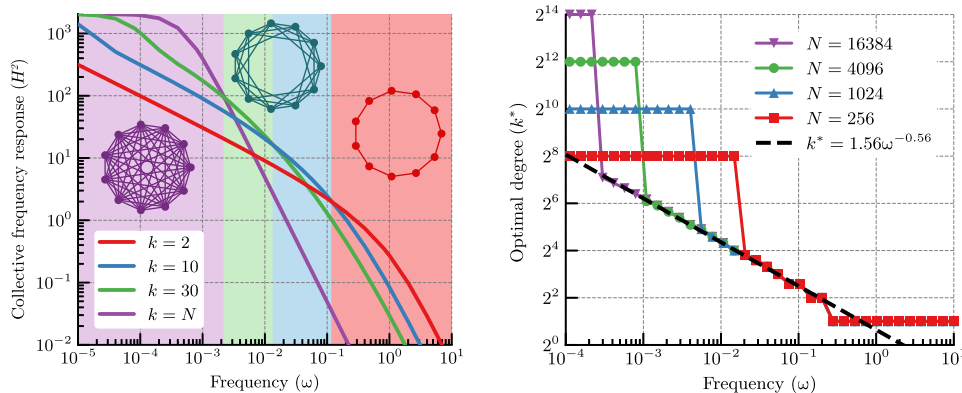


Fig. 1. Collective frequency response for a ring network. (Left) Response of $N = 2048$ agents performing distributed linear consensus over a regular periodic 1D grid with fixed degree k . Larger degrees yield a higher response at low frequency ($\omega < 2 \times 10^{-3}$), while the opposite is true at high frequency ($\omega > \omega_{\text{high}} = 0.278$). (Right) Optimal degree k^* for maximum collective response as a function of the frequency ω for a system of N agents distributed on a ring. For low frequency, the optimal k^* corresponds to an all-to-all connectivity. At higher frequencies, $k^*(\omega)$ follows the “bulk” behavior from Eq. 1 (fit, black line) up to its lowest possible value, $k^* = 2$, at $\omega = \omega_{\text{high}}$.

In general, the optimal degree for a given frequency depends on the network model considered. Figure 2 presents a comparison of the k^* values obtained for different kinds of networks. All networks considered display a k^* that decreases monotonically with frequency, transitioning from an all-to-all connectivity at low frequency to a minimal connectivity at high ones. Note that the minimum possible degree is $k = 2$ for all models except for the 2D mesh that requires a minimum of $k = 4$. This transition from all-to-all to minimal connectivity may be abrupt (2D mesh, random), continuous (caveman model, where $k^* \approx 1/\omega$), or a combination of both (1D ring).

Network optimization: Weights and structure

The results in the previous section present a clear phenomenology for several types of networks that hint toward a possible universal behavior. To explore how general this phenomenology is, we consider two cases for which we relax some of the assumptions made previously. In the first case, we consider weighted networks, for which the connection between agents is not binary. In the second case, we consider a general network optimization problem for a small system where no particular structure is imposed.

Optimal weights

Unweighted graphs, i.e., networks where $a_{ij} = \{0, 1\}$, represent only a small subset of the possible connectivities between agents, and there is a priori no reason to assume that the optimal connectivity can be represented by an unweighted graph. Thus, expanding the study to weighted graphs opens the possibility of finding networks with arbitrary distribution of weights that have a higher frequency response.

Using a linear parametrization of \mathbf{W} (see Materials and Methods) for a regular ring, where w_{ij} depends only on the topological distance between agents $d_{ij} = |i - j|$, the numerical optimization of Eq. 13 yields the connectivity profiles shown in Fig. 2. For a ring of $N = 4096$ agents, the optimal responsiveness at frequencies $\omega < 5 \times 10^{-3}$ corresponds to a simple all-to-all connectivity (all $w_{ij} = 1/N$). At higher frequencies, one obtains a smooth profile where the optimal connectivity decreases with distance. Still, most nodes are either connected ($w_{ij} \approx 1/k^*$) or disconnected ($w_{ij} \approx 0$), with only a few near the transition having intermediate weights. This profile is very similar to the unweighted case, and the effective number of neighbors inferred from the profile (see inset of Fig. 2) is identical to the values presented in the previous section.

This numerical optimization corroborates that allowing weighted connections does not sensibly change the phenomenology observed, where limiting the number of connections to a certain frequency-dependent value optimizes the response of the system.

Optimal structure

So far, we have considered particular network models for which the degree k can be controlled. All these models exhibit an optimal frequency response when the degree is set to a certain k^* that decreases with increasing frequency. However, the particular value of k^* depends on the model. The results do not guarantee that, in general, any network with a given fixed degree k^* will yield a higher collective response than any other network with a different degree.

Does the relationship between the number of neighbors and frequency response still hold when considering arbitrary network structures? Preliminary results of numerical discrete optimization for small systems suggest that it is the case. The bottom row of Fig. 2 shows examples of the networks obtained by performing simulated annealing optimization for the frequency response of a system of $N + 1 = 11$ agents for six frequencies ω .

At frequencies $\omega \leq 0.1$, the optimization procedure consistently yields an all-to-all connectivity. For higher frequencies $\omega = 0.2$ and 0.3 , the stochastic nature of the optimization procedure generates slightly different configurations, but all of them contain a clique of seven agents at $\omega = 0.2$ (5 to 6 at $\omega = 0.3$), with the rest of the agents in the periphery of the clique. These configurations have an average degree $\langle k \rangle = 4.7$ for $\omega = 0.2$ and $\langle k \rangle = 3.6$ for 0.3 . Further increasing the frequency yields disconnected graphs, where the agents are distributed in clusters of either four, three, or two agents for $\omega = 0.4, 0.6$, and 1 , respectively.

While the current results are limited to small systems, they show that the existence of an optimal mean degree that decreases with increasing frequency is not a particular feature of the network models considered here, but a general phenomenology of systems performing linear distributed consensus.

Application to swarm robotics

To showcase the application of these results to the design of interaction networks in the nascent field of swarm robotics, we perform a series of experiments on heading consensus with a swarm of land

Downloaded from <http://advances.sciencemag.org/> on June 15, 2019

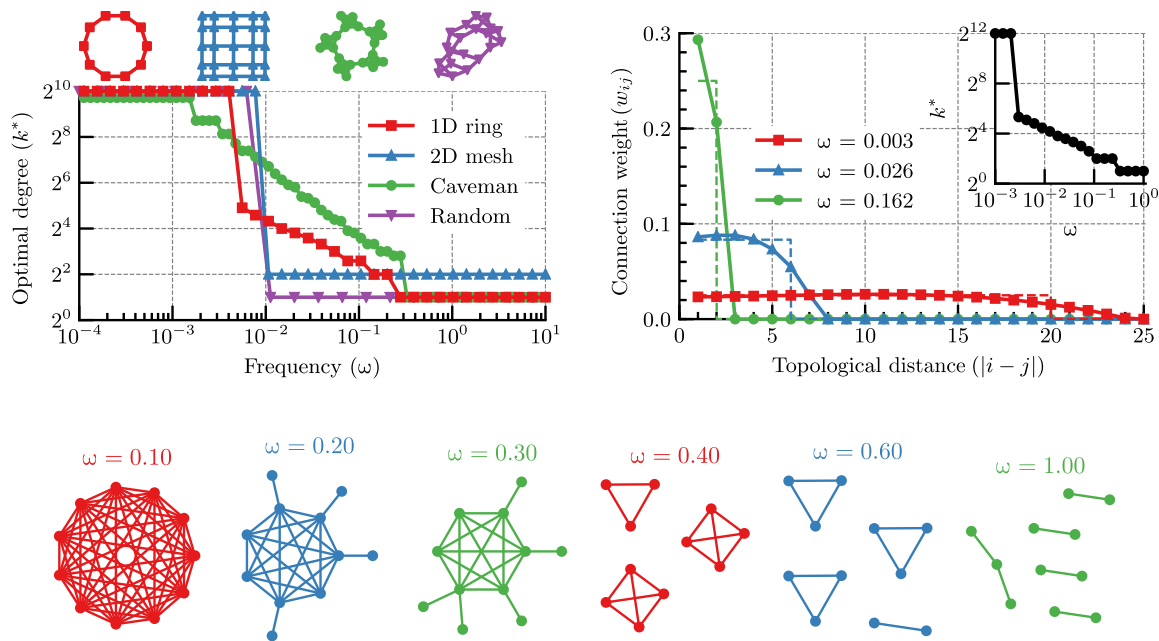


Fig. 2. Optimal networks for collective response. (Top left) Optimal degree k^* for maximum collective response as a function of the frequency ω for a system of $N = 1024$ arranged in different network topologies ($N = 840$ for the caveman topology). Note that some networks display a sudden transition from all-to-all to minimal connectivity, while others have an intermediate range of frequencies where k^* follows a scaling law of the form Eq. 1. (Top right) Optimal connection weights w_{ij} as a function of topological distance for a system of $N = 4096$ agents at a given frequency ω . The inset shows the optimal number of connections k^* obtained by fitting a Heaviside function to the weight distributions. (Bottom) Optimal network topologies for a system of $N + 1 = 11$ agents obtained by stochastic numerical discrete optimization over the space of unweighted, undirected graphs. Instead of fixing the leader to be a particular agent, the optimization maximizes the collective frequency response averaged over all the possible leaders (allowing disconnected graphs to be optimal). Note that the mean degree of these networks is consistently reduced with increased frequency ω .

robots where each one aligns its direction of motion with that of their neighbors. This form of consensus, inspired by Vicsek’s model of collective motion in natural swarms (30), is closely related to the first-order distributed linear consensus discussed previously. However, an empirical implementation involves significant deviations from the ideal scenario.

For instance, while the dynamics of the robots are ultimately governed by physical processes that are continuous in time, the robots sense each other’s state using asynchronous, discrete communications with stochastic delays. The consensus protocol itself, like in the Vicsek model, is nonlinear due to the state variable being an angular quantity as opposed to a scalar one.

The swarm is composed of 11 robots (see Materials and Methods) equipped with a custom “swarm-enabling” unit—providing data-processing power and distributed communications—that allows the robots to perform complex, decentralized cooperative control algorithms (14). Each robot periodically sends its heading to its k neighbors. One of the robots, the leader, is set to continuously rotate at a fixed frequency ω in the range $0.01 \text{ Hz} < \omega < 0.1 \text{ Hz}$. The rest of $N = 10$ robots are set to perform the heading consensus algorithm of Eq. 16. The experiments are run for at least four periods of leader rotation on an open space of around 20 m^2 . The experimental setup does not impose any significant restriction on the location, other than a relatively flat and regular ground and enough space for the robots to move freely.

The effect of connectivity on responsiveness is investigated by repeating the experiment for different values of the number of neighbors k and different rotation frequencies ω for the leader. See the Supplementary Materials for videos of this experiment performed at low and high

frequencies. A characteristic selection of the results obtained is presented in Fig. 3. When the robots are connected by means of a ring topology ($k = 2$, left column), the qualitative behavior of the swarm is not critically affected by the frequency: both at $\omega = 0.04 \text{ Hz}$ and 0.06 Hz , the four or five robots topologically closest to the leader are able to reasonably follow it, with the rest lagging behind.

However, the picture changes when an all-to-all connectivity underpins the operations of the swarm ($k = 10$, right column). Because each robot has access to the same information, they all behave identically. This not only allows the whole swarm to follow closely the leader at frequencies below a threshold of $\omega_c \approx 0.05 \text{ Hz}$ but also causes the response of the system to drop drastically above this threshold. In the context of the Vicsek model, the swarm has a low polarization when $k = 2$ (0.62 to 0.77) and a high one when $k = 10$ (0.91 to 0.95). Here, we see that a high polarization not only allows the system to have a large coordinated response at low frequency but also tends to “ossify” the collective, thus drastically reducing its response at high frequency.

This phenomenology is reminiscent of what has been observed in natural systems. Specifically, it is known that while flocks of starlings have high levels of ordering or polarization (31), swarms of midges display low levels of polarization but a large connected correlation (5), and thus high responsiveness to biologically relevant environmental perturbations (8).

The measured capacity of the swarm to align to the time-dependent orientation of a leader is presented in Fig. 4 as a function of the number of neighbors for a selected number of frequencies. These results confirm the predictions from the distributed linear consensus model applied to

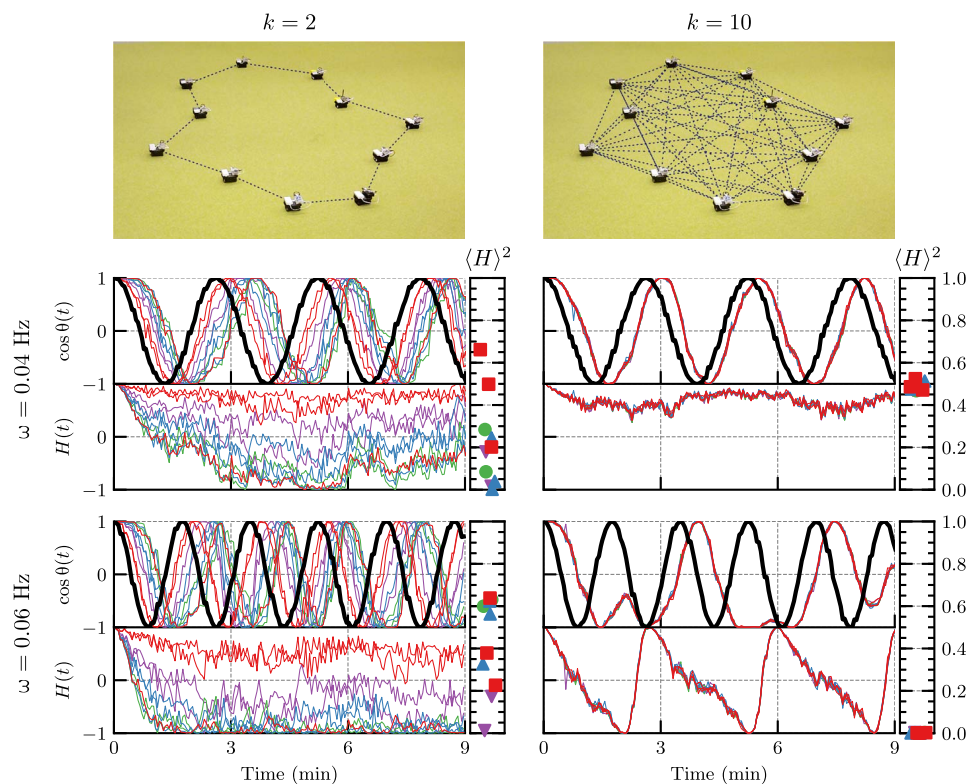


Fig. 3. Distributed heading consensus experiment with a swarm of 11 robots. Evolution of the robots' heading in an experiment with one leader (black line) rotating at frequency $\omega = 0.04$ Hz (top) or $\omega = 0.06$ Hz (bottom) when each robot has either $k = 2$ neighbors (left column) or $k = 10$ (right column). The degree by which each agent is following the leader at a given instant is $H(t)$ (Eq. 18), and the square of the mean value (displayed in the lateral bar) is its frequency response.

complex, realistic cases of collective motion. Specifically, the robots benefit from having more connections at low frequency but it hinders their responsiveness at high ones.

Unfortunately, our experimental system is too small to observe an optimal connectivity k^* different than two or N , or to be able to meaningfully study different network models. However, with the advent of large swarm robotics systems (32), it is likely that different values of k^* can be experimentally measured for various network topologies.

DISCUSSION

This paper presents a detailed study on the relation between the collective frequency response to a time-dependent, localized signal (a leader) in a multi-agent system performing distributed consensus and the connectivity between the agents. The study shows that the response to the driving signal arising from different connectivities depends critically on the time scale of the signal: The smaller the time scale, the better low-degree connectivities perform as compared with high-degree ones.

In particular, we observe that when the agents are connected by means of a static k -nearest neighbor ring configuration, the collective response of the system—excluding finite-size effects at low frequency—is maximized for a particular frequency-dependent number of neighbors that follows a scaling law of the form $k^* \propto \omega^{-0.56}$. This functional form for $k^*(\omega)$ is not universal, but all the network models considered here consistently display a k^* monotonically decreasing with the signal frequency ω . For instance, a caveman network has a $k^* \propto \omega^{-1}$,

and other networks such as a bidimensional grid or a regular random network have a bimodal k^* that corresponds to an all-to-all connectivity for $\omega < \omega_c$ and to a minimal connectivity for $\omega > \omega_c$.

We posit that the existence of this optimal connectivity k^* , dependent on the time scale $\tau = 1/\omega$ of the local perturbation, has eluded prior investigations on distributed consensus because those did not consider either the localized nature of perturbations (4, 6, 16, 18, 21, 23, 29, 30) or the short-time transient component of the response (4, 7, 13, 16, 26). On the one hand, the local injection of the perturbation into the collective dynamics is key to the uncovered phenomenon because it leads to a complex propagation of the input signal through a host of feedback and feedforward loops associated with the topology of the network. On the other hand, the short-time (high frequency) transient responsive behavior is not captured if only steady-state deviations from consensus, integrated metrics, such as consensus speed, or similar proxies are used to characterize the response.

This phenomenology stands in stark contrast with that of the well-studied consensus-reaching process: The prescription for a fast consensus is not necessarily compatible with the one for a responsive consensus. The speed of consensus—also known as convergence rate—can readily be obtained from the second-smallest eigenvalue λ_2 of the Laplacian matrix in the absence of a leader (24), or from the smallest eigenvalue λ_1^g of the grounded Laplacian matrix in the presence of a leader (25). In general, λ_2 and λ_1^g tend to increase with increasing mean degree (18, 25, 28). In particular, the extreme case of an all-to-all connectivity ($k = N$) maximizes the speed of consensus. This speed defines the response's cutoff frequency, and thus, an increase in it improves the collective response to slow-changing environmental

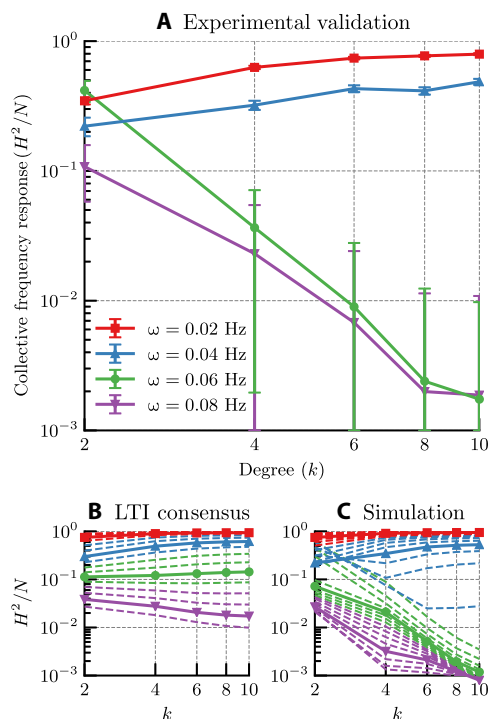


Fig. 4. Collective response in leader-follower heading consensus for a system of $N + 1 = 11$ agents. (A) Experimental collective frequency response (Eq. 19, normalized) obtained with 10 robots performing distributed heading consensus plus one leader rotating at a fixed frequency ω . (B) Response for the equivalent LTI distributed linear consensus (Eq. 5, normalized) with the same N . (C) Response obtained with simulations of the heading consensus algorithm (Eq. 16).

perturbations. However, the short time-scale response of a decentralized system (i.e., subjected to higher frequency components) cannot be inferred solely from the large time scale one corresponding to the consensus speed. In that regime, the response is found to have a nontrivial relationship with the degree of the interaction network. All the network models presented here show an intrinsic trade-off in this respect: Any change that increases the low-frequency response decreases the high-frequency one, and vice versa.

The linear distributed consensus is a general model with applications in a wide range of social and artificial systems. In general, these systems will have different design constraints in connectivity (arising, for example, from a spatial embedding of the nodes), which make certain network models more relevant than others. We performed two exploratory analyses that indicate that these results are relatively general. On the one hand, when the ring network is allowed to have weighted links, the optimal distribution of weights is comparable to the distribution in the unweighted case. On the other hand, when no structure is imposed, discrete optimization on small systems shows that the optimal degree distribution is still such that $\langle k \rangle^*$ decreases with ω .

In practice, many multi-agent applications—such as artificial swarms—require an explicit or implicit design of the interaction network or interaction model. This process usually involves numerous trade-offs due to numerous design constraints (2), and thus, the network is designed to facilitate certain desired qualities of the system—such as robustness, scalability, and responsiveness. What the findings presented here teach us is that, for multi-agent systems to be

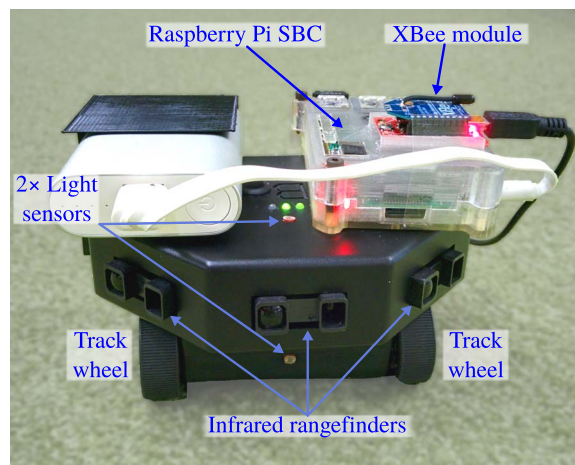


Fig. 5. Robotic platform used in the experiments. The SBC and XBee module attached to the robot provides autonomy and distributed communications to the unit, making it able to swarm and perform decentralized collective motion such as heading consensus. (Photo credit: David Mateo, Singapore University of Technology and Design.)

responsive, there is not one optimal interaction network that can guarantee responsiveness in general. Instead, the connectivity can only be optimized to yield maximal responsiveness at a particular time scale. For all the cases considered here, an increase in responsiveness to low-frequency signals implies a decrease in the response to high frequencies, and vice versa.

To showcase the applicability of our findings to the practical arena, we have performed a series of experiments with a set of 11 swarming robots that seek to emulate the behavior of a leader by performing a classical example of collective motion. These robots communicate their state to a given set of neighbors and align their direction of motion to the average of said neighbors. The experiments confirm that, when the leader changes its direction slowly, the agents are better at following the leader the more connected they are and that the opposite is true for fast changes.

In this work, we have considered how the degree of the network affects the responsiveness of multi-agent systems performing distributed consensus. However, it is not possible to change the degree distribution of a network without also changing its other properties such as clustering or shortest path. While it is known that at high enough frequency the response is determined only by the mean degree (8), the question of which structural properties of the network are good predictors of a system’s response at different regimes remains open.

MATERIALS AND METHODS

Distributed linear consensus

Let us consider a group of $N + 1$ identical agents performing a distributed consensus protocol on their scalar state-variable $x_i(t)$. The dynamics of the system is determined by the state vector $\mathbf{X}(t) = \{x_i(t); i = 0, \dots, N\}$ and the adjacency matrix of the underlying graph $\mathbf{A} = \{a_{ij}; i, j = 0, \dots, N\}$, where $a_{ij} = 1$ if agent i is connected to j and 0 otherwise. Given a certain connectivity graph, the system evolves according to

$$\frac{dx_i}{dt} = \frac{\omega_0}{k_i} \sum_{j=0}^N a_{ij}(x_j(t) - x_i(t)) = \sum_{j=0}^N w_{ij}x_j(t) \tag{2}$$

where ω_0 is the natural response frequency of our identical agents, and $k_i = \sum_{j=0}^N a_{ij}$ is the degree (or number of neighbors) of agent i . The quantity $w_{ij} = \omega_0(a_{ij}/k_i - \delta_{ij})$ —where δ_{ij} is a Kronecker delta—is introduced for the sake of a compact notation. Note that, by definition, $w_{ii} = -\omega_0$ and $\sum_j w_{ij} = 0$ for all i .

To model the response of the system, we consider a leader-follower consensus scenario where one agent—for example, agent $i = 0$, the leader—does not abide by the dynamics of Eq. 2 but instead follows an arbitrary trajectory $x_0(t) = u(t)$. In the presence of this single leader, Eq. 2 can be recast as

$$\frac{dx_i}{dt} = \sum_{j=1}^N w_{ij}x_j(t) + w_{i0}u(t) \tag{3}$$

for $i = 1, \dots, N$. The solution of Eq. 3 (up to an integration constant) can be written compactly in matrix notation on the frequency domain as

$$\mathbf{X}(\omega) = (i\omega\mathbf{I} - \mathbf{W}_F)^{-1}\mathbf{W}_L u(\omega) \tag{4}$$

where \mathbf{I} is the identity matrix of dimension N , $\mathbf{W}_F = \{w_{ij}\}$ is the $N \times N$ consensus protocol matrix between the follower agents (also known as state matrix \mathbf{A} in LTI systems), and $\mathbf{W}_L = \{w_{i0}\}$ is the $N \times 1$ consensus protocol matrix between the followers and the leader (also known as input matrix \mathbf{B}).

The response function or susceptibility measures the capacity of the multi-agent system to follow the leader’s trajectory, $u(t)$, and can be expressed in the frequency domain (33) as

$$\mathbf{H}(\omega) = \left(\frac{\delta\mathbf{X}}{\delta u}\right)(\omega) = (i\omega\mathbf{I} - \mathbf{W}_F)^{-1}\mathbf{W}_L \tag{5}$$

The entries of the vector $\mathbf{H} = \{h_i\}_{i=1,\dots,N}$ correspond to the frequency response of each individual agent, with $|h_i(\omega)| \leq 1$ for all i and ω (33). As is clear from Eq. 5, the response functions have a nontrivial dependency on the topology of the agents’ connectivity through the entries of \mathbf{W}_F and \mathbf{W}_L . The collective response of the system can be characterized by performing a singular value decomposition of \mathbf{H} , giving a single singular value $\sigma^2 = \sum_i |h_i|^2 = H^2$. Throughout this work, we will use $H^2(\omega)$ (or its normalized form, H^2/N) as a measure of the collective frequency response of multi-agent systems. From Eq. 4, one can see that if $u(t)$ is taken to be a white-noise stochastic perturbation, the fluctuations in the collective will be correlated as $\langle \mathbf{X}^\dagger \mathbf{X} \rangle \propto H^2$.

Models of interaction network

To study the effect of the number of neighbors on the collective frequency response, we have considered a set of network topologies including 1D and 2D regular periodic lattices (i.e., a ring and a mesh), connected caveman model graphs (34), and regular random graphs.

We have chosen these models because they feature a fairly regular structure—minimizing the effect of the leader’s location—and they allow a fine control of the degree of the agents from $k = 2$ or 4 up to an all-to-all connectivity ($k = N$). They are also guaranteed to be connected and thus have $H^2(\omega = 0) = N$ for any k , with the exception of the

random regular graph at low degrees ($k \lesssim \log N$) (35). All these models provide unweighted, undirected interaction networks without self-loops.

Regular lattice

These networks are constructed by placing the agents in a regular square grid embedded in an Euclidean d -dimensional periodic space (a ring for $d = 1$ and a torus for $d = 2$) and connecting each node to its k nearest neighbors on the grid. This structure guarantees that all nodes have the same centrality and degree, and we can set the leader arbitrarily as the first agent without loss of generality.

For the 1D ring, the degree can be any even value between 2 and N . For the 2D mesh, only multiples of 4 are valid degrees.

Caveman model

The connected caveman graph is composed of n complete subgraphs with $k + 1$ nodes, where one edge from each cluster is modified so that it connects neighboring clusters (34). Note that the number of nodes is $N = (k + 1)n$ and the average degree is k (there are n nodes with a degree of $k + 1$ and another k with $k - 1$; the rest have a degree of k).

We choose the number of agents N to be a highly composite number (such as 720, 840, or 1260) to have a high resolution on the possible values of the degree k while keeping N constant.

Regular random networks

Last, we consider the case of regular random graphs, i.e., graphs randomly sampled from all the possible k -regular graphs with N nodes. The graphs sampled are expected to be less structured than the lattices or caveman graphs, typically having significantly lower diameter and clustering coefficient. The frequency response of random networks presented in Fig. 2 is obtained by averaging over 100 samples for each frequency.

Weight optimization

The optimal distribution of weights \mathbf{A} (or equivalently \mathbf{W}) for collective frequency response can be obtained by solving

$$\frac{\delta H^2}{\delta \mathbf{W}} = 0 \tag{6}$$

with the conditions that $\sum_{j=0}^N w_{ij} = 0$ for all j . This gradient can be written as

$$\begin{aligned} \frac{\partial H^2}{\partial w_{ij}} &= [\mathbf{W}_L^\dagger (i\omega - \mathbf{W}_F)^{-1\dagger} (i\omega - \mathbf{W}_F)^{-2}]_i (\mathbf{W}_L)_j + \text{h.c.} \\ \frac{\partial H^2}{\partial w_{i0}} &= [\mathbf{W}_L^\dagger (i\omega - \mathbf{W}_F)^{-1\dagger} (i\omega - \mathbf{W}_F)^{-1}]_i + \text{h.c.} \\ \frac{\partial H^2}{\partial w_{0j}} &= 0 \end{aligned} \tag{7}$$

where h.c. stands for the Hermitian conjugate of the preceding expression.

If no constraints are imposed on the weights, a trivial solution is obtained in which the response is maximized by having all the agents connected only to the leader. For this study to be applicable to cases where the leader is not known in advance and where it may even change with time, one needs to impose some additional symmetries either in H^2 or in \mathbf{W} directly.

One option to introduce the symmetry in H^2 is to compute the frequency response not by fixing the leader to be a particular agent but to

average over all the possible leaders instead. This option is appropriate for small systems, but for relatively large systems, it is more suitable to impose symmetries on \mathbf{W} instead to reduce the number of variables (which otherwise grows as N^2).

To impose arbitrary symmetries in the system, one can constrain the weights with a given parametrization

$$w_{ij} = F(i, j, \{c_k\}) \quad (8)$$

where $\{c_k\}$ is a set of free parameters. The gradient of H^2 with respect to these parameters is

$$\begin{aligned} \frac{\partial H^2}{\partial c_k} = & \sum_{ij} \frac{\partial H^2}{\partial w_{ij}} \frac{\partial w_{ij}}{\partial c_k} \mathbf{W}_L^\dagger (i\omega - \mathbf{W}_F)^{-1\dagger} (i\omega - \mathbf{W}_F)^{-2} \frac{\partial \mathbf{W}_F}{\partial c_k} \mathbf{W}_L + \\ & + \mathbf{W}_L^\dagger (i\omega - \mathbf{W}_F)^{-1\dagger} (i\omega - \mathbf{W}_F)^{-1} \frac{\partial \mathbf{W}_L}{\partial c_k} + \text{h.c.} \end{aligned} \quad (9)$$

We consider the case where the connection between agents depends only on the topological distance between them. Given a measure of topological distance $d(i, j)$, this condition can be written as a linear parametrization of the form

$$w_{ij} = \sum_k c_k m_{ij}^k \quad (10)$$

where

$$m_{ij}^k = \begin{cases} 1 & \text{if } d(i, j) = k \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

With a linear parametrization, we obtain a close form for the gradient as

$$\begin{aligned} \frac{\partial H^2}{\partial c_k} = & \mathbf{W}_L^\dagger (i\omega - \mathbf{W}_F)^{-1\dagger} (i\omega - \mathbf{W}_F)^{-2} \mathbf{M}^k \mathbf{W}_L + \\ & + \mathbf{W}_L^\dagger (i\omega - \mathbf{W}_F)^{-1\dagger} (i\omega - \mathbf{W}_F)^{-1} \mathbf{M}_L^k + \text{h.c.} \end{aligned} \quad (12)$$

where $\mathbf{M}^k = \{m_{ij}^k\}$ and $\mathbf{M}_L^k = \{m_{i0}^k\}$.

The normalization of the weights during optimization can be imposed through a Lagrange multiplier. Thus, instead of maximizing H^2 , we define a cost function of the form

$$\mathcal{L} = H^2 - \frac{\lambda}{2} \sum_{i=0}^N \left| \sum_{j=0}^N w_{ij} \right|^2 \quad (13)$$

Using this cost function, the optimization problem over w_{ij} with $i \neq j$ (the diagonal terms are fixed to $w_{ii} = -1$) can be written as

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial w_{ij}} = & \frac{\partial H^2}{\partial w_{ij}} - \lambda \sum_{l=0}^N w_{il} = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} = & \frac{1}{2} - \sum_{ij} (\mathbf{W}^\dagger \mathbf{W})_{ij} = 0 \end{aligned} \quad (14)$$

For the linear parametrization of Eq. 10, the gradient of the cost function with respect to c_k can be written as

$$\frac{\partial \mathcal{L}}{\partial c_k} = \frac{\partial H^2}{\partial c_k} - \lambda \sum_{ij} (\mathbf{W}^T \mathbf{M}^k)_{ij} \quad (15)$$

Robotic platform

For the experimental validation, we have used a differential-drive robot developed in-house as a low-cost tool for robotics research (36). The platform is equipped with six infrared rangefinders, an inertial measurement unit (IMU), two wheel encoders, and two light sensors.

The robots are granted collective autonomy and distributed communication by attaching a “swarm-enabling” unit (14) composed of a single-board computer (SBC) and an XBee module (see Fig. 5) (37). This unit interfaces with the robot via Bluetooth and is responsible for implementing the cooperative control algorithms, communicating with other robots, and controlling the motion of the robot.

The units communicate via radio signals sent over the distributed, dynamic mesh established by the XBee modules. Because physical limitations impose a maximum range at which these signals can be sent, the swarm has a natural “metric” interaction model, meaning that each agent is able to communicate with any other agents within a given distance. Field experiments with aquatic autonomous surface vehicles using a similar setup for communication (15) have shown that, when the number of agents is large, the interaction model is also weakly density-dependent, deviating from a pure metric model. However, for the current experimental setup, there is no practical spatial limitation on communication between the agents, and instead, the different interaction networks are explored by tuning the cooperative control rule.

To control the robot, the SBC processes the data from the IMU and encoders through a Kálmán filter to have an accurate estimation of the platform’s location. Then, a proportional integral derivative (PID) controller allows it to adjust the trajectory of the robot. The PID coefficients are tuned using the Ziegler-Nichols frequency response method (33).

The behavior of the robot, defined by the cooperative control algorithm, is implemented in the swarm-enabling unit using the in-house marabunta package (14). The software is designed prioritizing platform-agnostic development of cooperative control rules, portability, and a simple workflow to facilitate fast prototyping. In the study of collective motion, these behavior rules typically take the form of simple, local, iterative algorithms, where the velocity of an agent is defined in terms of its own state, the local environment, and the state of the agents in a certain neighborhood.

Heading consensus experiments

To measure empirically the collective frequency response of a robotic swarm, we have performed leader-follower heading consensus (38) experiments using 11 robots. One of them is designated as the leader and its behavior is a simple rotational motion at constant frequency ω . We label it as leader because its behavior does not depend on the rest of the swarm, but the agents have no way of distinguishing the leader from any other agent. The other 10 “follower” robots perform a heading consensus algorithm using the information received from a given set of neighboring agents, irrespective of whether a neighbor is a leader or a follower.

The heading consensus algorithm determines a target heading $\bar{\theta}$ for each robot of the form

$$\bar{\theta}_i = \langle \theta_j \rangle_{j \sim i} = \arctan \left(\frac{\sum_{j \sim i} \sin \theta_j}{\sum_{j \sim i} \cos \theta_j} \right) \quad (16)$$

where $j \sim i$ denotes the neighbors of i and $\langle \cdot \rangle$ an angular average. Each follower robot updates its target heading asynchronously every $\Delta T = 0.1$ s. The information of each neighbor's state is also updated every 0.1 s, but not necessarily concurrently with each other or with the update of Eq. 16. The robots are interconnected with a ring network topology such that each individual unit can only communicate with k other robots.

This heading consensus algorithm used in the experiments is a discrete-time equivalent to a linear consensus algorithm where the modulus of the state variable is fixed, i.e., it is equivalent to

$$\frac{d\theta_i}{dt} = \omega_0 \left(\langle \theta_j \rangle_{j \sim i} - \theta_i \right) \quad (17)$$

if $\omega_0 \Delta T \gg 1$ and for times $t \gg \Delta T$.

Because the consensus algorithm is nonlinear, we cannot use the LTI framework to define the frequency response as Eq. 5. Instead, we use an equivalent response metric by defining the state of an agent as $x_i(t) = e^{-i\theta_i(t)}$, where θ_i is its heading. The capacity of an agent i to follow the leader L is then measured by

$$H_i = \text{Re} \left[\frac{1}{T} \int_0^T \frac{x_i(t)}{x_L(t)} dt \right] = \frac{1}{T} \int_0^T \cos(\theta_i - \theta_L) dt \quad (18)$$

where T is the duration of the experiment and $|H_i| \leq 1$. The collective frequency response of the system is then measured by

$$H^2 = \sum_{i=1}^N |H_i|^2 \quad (19)$$

with $0 \leq H^2 \leq N$.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/5/4/eaau0999/DC1>

Movie S1. Experiment to measure the collective response in leader-follower heading consensus of a swarm of $N + 1 = 11$ land robots with a low-frequency input signal.
 Movie S2. Experiment to measure the collective response in leader-follower heading consensus of a swarm of $N + 1 = 11$ land robots with a high-frequency input signal.

REFERENCES AND NOTES

1. J. Krause, G. D. Ruxton, *Living in Groups* (Oxford Series in Ecology and Evolution, Oxford Univ. Press, 2002).
2. R. Bouffanais, *Design and Control of Swarm Dynamics* (Springer, 2016).
3. A. Cavagna, I. Giardina, T. S. Grigera, The physics of flocking: Correlation as a compass from experiments to theory. *Phys. Rep.* **728**, 1–62 (2018).
4. G. F. Young, L. Scardovi, A. Cavagna, I. Giardina, N. E. Leonard, Starling flock networks manage uncertainty in consensus at low cost. *PLOS Comput. Biol.* **9**, e1002894 (2013).
5. A. Attanasi, A. Cavagna, L. Del Castello, I. Giardina, S. Mellillo, L. Parisi, O. Pohl, B. Rossaro, E. Shen, E. Silvestri, M. Viale, Collective behaviour without collective order in wild swarms of midges. *PLOS Comput. Biol.* **10**, e1003697 (2014).

6. T. Vicsek, A. Zafeiris, Collective motion. *Phys. Rep.* **517**, 71–140 (2012).
7. D. T. Swain, I. D. Couzin, N. E. Leonard, Coordinated speed oscillations in schooling killifish enrich social communication. *J. Nonlinear Sci.* **25**, 1077–1109 (2015).
8. D. Mateo, Y. K. Kuan, R. Bouffanais, Effect of correlations in swarms on collective response. *Sci. Rep.* **7**, 10388 (2017).
9. Y. Shang, R. Bouffanais, Consensus reaching in swarms ruled by a hybrid metric-topological distance. *Eur. Phys. J. B* **87**, 294 (2014).
10. J. H. Fowler, N. A. Christakis, Cooperative behavior cascades in human social networks. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 5334–5338 (2010).
11. D. Centola, The spread of behavior in an online social network experiment. *Science* **329**, 1194–1197 (2010).
12. V. Amelkin, F. Bullo, A. K. Singh, Polar opinion dynamics in social networks. *IEEE Trans. Automat. Contr.* **62**, 5650–5665 (2017).
13. W. Ren, R. W. Beard, *Distributed Consensus in Multi-Vehicle Cooperative Control: Theory and Applications, Communications and Control Engineering Series* (Springer, 2008).
14. M. Chamanbaz, D. Mateo, B. M. Zoss, G. Tokić, E. Wilhelm, R. Bouffanais, D. K. P. Yue, Swarm-enabling technology for multi-robot systems. *Front. Robot. AI* **4**, 12 (2017).
15. B. M. Zoss, D. Mateo, Y. Kong Kuan, G. Tokić, M. Chamanbaz, L. Goh, F. Vallega, R. Bouffanais, D. K. P. Yue, Distributed system of autonomous buoys for scalable deployment and monitoring of large waterbodies. *Auton. Robots* **42**, 1669–1689 (2018).
16. F. Lin, M. Fardad, M. R. Jovanović, Algorithms for leader selection in stochastically forced consensus networks. *IEEE Trans. Automat. Contr.* **59**, 1789–1802 (2014).
17. M. Komareji, R. Bouffanais, Resilience and controllability of dynamic collective behaviors. *PLOS ONE* **8**, e82578 (2013).
18. Y. Shang, R. Bouffanais, Influence of the number of topologically interacting neighbors on swarm dynamics. *Sci. Rep.* **4**, 4184 (2014).
19. G. Punzo, G. F. Young, M. Macdonald, N. E. Leonard, Using network dynamical influence to drive consensus. *Sci. Rep.* **6**, 26318 (2016).
20. Y.-Y. Liu, A.-L. Barabási, Control principles of complex systems. *Rev. Mod. Phys.* **88**, 035006 (2016).
21. M. Komareji, Y. Shang, R. Bouffanais, Consensus in topologically interacting swarms under communication constraints and time-delays. *Nonlinear Dyn.* **93**, 1287–1300 (2018).
22. A. Sekunda, M. Komareji, R. Bouffanais, Interplay between signaling network design and swarm dynamics. *Network Science* **4**, 244–265 (2016).
23. Y. Cao, W. Yu, W. Ren, G. Chen, An overview of recent progress in the study of distributed multi-agent coordination. *IEEE Trans. Industr. Inform.* **9**, 427–438 (2013).
24. R. Olfati-Saber, J. A. Fax, R. M. Murray, Consensus and cooperation in networked multi-agent systems. *Proc. IEEE* **95**, 215–233 (2007).
25. M. Pirani, S. Sundaram, On the smallest eigenvalue of grounded Laplacian matrices. *IEEE Trans. Automat. Contr.* **61**, 509–514 (2016).
26. A. Clark, B. Alomair, L. Bushnell, R. Poovendran, Leader selection in multi-agent systems for smooth convergence via fast mixing, in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)* (IEEE, 2012), pp. 818–824.
27. Z. Li, Z. Duan, G. Chen, L. Huang, Consensus of multiagent systems and synchronization of complex networks: A unified viewpoint. *IEEE Trans. Circuits Syst. I* **57**, 213–224 (2010).
28. J. A. Almendral, A. Diaz-Guilera, Dynamical and spectral properties of complex networks. *New J. Phys.* **9**, 187 (2007).
29. H. Chaté, F. Ginelli, G. Grégoire, F. Raynaud, Collective motion of self-propelled particles interacting without cohesion. *Phys. Rev. E* **77**, 046113 (2008).
30. T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, O. Shochet, Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.* **75**, 1226–1229 (1995).
31. A. Attanasi, A. Cavagna, L. Del Castello, I. Giardina, T. S. Grigera, A. Jelić, S. Mellillo, L. Parisi, O. Pohl, E. Shen, M. Viale, Superfluid transport of information in turning flocks of starlings. *Nat. Phys.* **10**, 691–696 (2014).
32. M. Rubenstein, A. Cornejo, R. Nagpal, Programmable self-assembly in a thousand-robot swarm. *Science* **345**, 795–799 (2014).
33. K. Ogata, *Modern Control Engineering* (Prentice Hall, ed. 5, 2010).
34. D. J. Watts, Networks, dynamics, and the small-world phenomenon. *Am. J. Sociol.* **105**, 493–527 (1999).
35. B. Bollobás, *Modern Graph Theory* (Springer, 1998), pp. 215–252.
36. T. G. Karimpanal, M. Chamanbaz, A. Gupta, W. Lizheng, T. Jeruzalski, E. Wilhelm, Adapting Low-Cost Platforms for Robotics Research. *Proceedings FinE-R workshop, The Path to Success: Failures in rEAL Robots, International Conference on Intelligent Robots and Systems (IROS 2015)*, Hamburg, Germany, 2 October 2015, pp. 20–26.
37. R. Piyare, S.-r. Lee, Performance analysis of XBee ZB module based wireless sensor networks. *Int. J. Sci. Eng. Res.* **4**, 1615–1621 (2013).
38. W. Ren, R. W. Beard, Consensus seeking in multiagent systems under dynamically changing interaction topologies. *IEEE Trans. Autom. Control* **50**, 655–661 (2005).

Acknowledgments

Funding: This work was supported by MOE-Tier 1 grant SUTDT2017001 and SUTD-MIT International Design Centre grant IDG31700107. **Author contributions:** D.M. and

R.B. designed the study and the experiment. D.M. and N.H. developed the analytical and numerical tools and analyzed the data. V.H. and M.C. performed the experiments. All authors wrote the manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be requested from the authors.

Submitted 13 May 2018
Accepted 11 February 2019
Published 3 April 2019
10.1126/sciadv.aau0999

Citation: D. Mateo, N. Horsevad, V. Hassani, M. Chamanbaz, R. Bouffanais, Optimal network topology for responsive collective behavior. *Sci. Adv.* **5**, eaau0999 (2019).

Optimal network topology for responsive collective behavior

David Mateo, Nikolaj Horsevad, Vahid Hassani, Mohammadreza Chamanbaz and Roland Bouffanais

Sci Adv 5 (4), eaau0999.

DOI: 10.1126/sciadv.aau0999

ARTICLE TOOLS

<http://advances.sciencemag.org/content/5/4/eaau0999>

SUPPLEMENTARY MATERIALS

<http://advances.sciencemag.org/content/suppl/2019/04/01/5.4.eaau0999.DC1>

REFERENCES

This article cites 31 articles, 3 of which you can access for free
<http://advances.sciencemag.org/content/5/4/eaau0999#BIBL>

PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. 2017 © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. The title *Science Advances* is a registered trademark of AAAS.