A Model of Flocking Using Sheaves

Joseph K. Geisz
Department of Mathematics
Colorado State University
Fort Collins, CO
joegeisz@colostate.edu

Abstract

Sheaves have been used recently to model information on networks, such as the spread of opinions in a social network. Dynamical systems on networks model the evolving states of nodes on graphs. Using these lenses of sheaf theory and network dynamics, we explore a model of flocking. We describe from this perspective what it means for birds to come to consensus on flight velocities, and a system of ordinary differential equations (ODEs) that describes this consensus process. Then we couple these consensus dynamics with flight dynamics to describe a model of flocking. We include numerous visualizations of examples in 2 dimensions.

1 Introduction

Flocking has been studied extensively as an example of an emergent phenomena: local interactions between birds (or other species, though the term "schooling" is generally used for fish) give rise to patterns and large-scale behavior. Numerous models of flocking have been proposed and studied, the most famous being the "boids" algorithm [16], often used in computer graphics.

The study of dynamical systems on networks asks questions about how properties of a network or graph affect the dynamics of the state of each node. For example, the generalized Lotka-Volterra equations model how complex food webs affect the population of interconnected species [11]. The famous Kuramoto model illustrates the phenomena of synchronization [13]. Even the process of training a neural network can be viewed as a dynamical system [2].

Dynamical systems on networks can model opinion consensus. In a sense flocking is a type of consensus. The birds can choose any direction in which to fly, but some are preferable based on what they know about their neighbors. Flocking can be framed as a dynamical system on a network. The network's nodes are birds, and the edges are other birds that they interact with. However, unlike in the examples mentioned above the movement of the birds can change the underlying network as birds move closer or further from their neighbors. This situation is often called an adaptive network - the network affects the dynamics and the dynamics affect the network [1].

In this paper we first describe traditional graph consensus, using a well-studied matrix called the graph Laplacian. Then we show how a more complex system - which is directly analogous to graph consensus - can model how a coherent velocity can be chosen for each bird in a flock, and what it means for them to be coherent. This model was described in [8], but we explore the behavior in more detail. It relies on a matrix called the sheaf Laplacian, a generalization of the graph Laplacian. This system however does not account for movement. A second model is introduced in which the velocities dynamically update while the birds' positions evolve as well, a fully adaptive model. Examples are illustrated throughout. Other more complex models of dynamic flocking have been proposed, such as in [15, 12, 17], but these rely on the traditional graph Laplacian or the Kronecker product of the graph Laplacian and an identity Matrix. Replacing this matrix with the sheaf Laplacian changes the dynamics.

1st Conference on Topology, Algebra, and Geometry in Data Science (TAG-DS 2025).

The formulation of these models rely on a mathematical construction called a sheaf, which organizes data over a graph (or higher order network). This construction comes originally from algebraic geometry, but in recent years has been used in applied settings. It is only described in the context of graphs in this paper, but is easily extended to higher-order networks. In this paper we construct a sheaf Laplacian, analogous to a graph Laplacian, and explore some of its properties and how it is used in our model of flocking. See [7] for more background on sheaf Laplacians.

2 Consensus on a graph

Birds in a flock must come to a consensus on coherent velocities in which to fly. As background, we describe an example of a dynamical system on a network which models opinion consensus. The dynamics are well known as "consenus" or "diffusion" on a graph. We consider a graph $G = (V, E)$ which models a social network. V is the set of vertices or nodes, representing the people in the social network. E is the set of edges or connections, representing social connections. Each individual, out of N individuals total, is represented by a node $v_i \in V$, and the edges (generically denoted as $(v_i, v_j) \in E$) are weighted with positive numbers w_{ij} representing the strength of the bond/friendship between individuals v_i and v_j. For two unconnected individuals, w_{ij} is zero. This data is stored in a **weighted adjacency matrix** A with entries $a_{ij} = w_{ij}$. Each individual has a current state/opinion, represented by a real number. We are interested in the evolution of this opinion as time progresses so we say individual v_i's opinion at time t is $x_i(t)$. A derivation is included in appendinx C, and are described with the ODE:

$$\frac{d}{dt}x_i(t) = -x_i \deg(v_i) + \sum_{j=1}^{N} w_{ij}x_j \tag{1}$$

Where $\deg(v_i) = \sum w_{ij}$, summing over j, the indices of node v_i's neighbors. This is a linear system of ODE's. The **graph Laplacian** is the matrix $L = D - A$ where D is the diagonal matrix with the $d_{ii} = \deg(v_i)$ and A is the weighted adjacency matrix. So L has elements

$$l_{ij} = \begin{cases} \deg(v_i) & \text{if } i = j \\ -w_{ij} & \text{if } (v_i, v_j) \in E \\ 0 & \text{otherwise} \end{cases}$$

It is interesting to note that there is an equivalent way to define L. We first define: The **boundary matrix** of a graph G is a $|V| \times |E|$ matrix D with elements

$$d_{ij} = \begin{cases} w_{i,j} & \text{if } v_i \text{ is first vertex in } e_j \\ -w_{ij} & \text{if } v_i \text{ is the second vertex in } e_j \\ 0 & \text{otherwise} \end{cases}$$

Notice this requires an ordering on the edges, equivalent to an orientation on the graph. While changing the ordering/orientation will change the boundary matrix, no matter how you orient the edges it will always be the case that $L = DD^T$. This is an equivalent definition of the graph Laplacian.

Defining L allows us to write out the system (1), with the entire state vector $\vec{x} = (x_i)$, as

$$\frac{d}{dt}\vec{x}(t) = -L\vec{x} \tag{2}$$

Thus we have a linear, homogeneous system of ordinary differential equations that represent the evolution of the opinions of each individual. The Laplacian matrix is positive semi-definite, and has a zero eigenvalue for each connected component of the underlying graph. A basis for the kernel of L consists of one vector per connected component of the graph. Each vector consists of ones on a single connected component and zeros on all other components. This implies that the dimension of the kernel is the number of connected components in the graph. The positive eigenvalues (along with the minus sign in the ODE) tell us that the equilibrium solutions are stable. Thus all solutions will converge to a steady state where everyone agrees in each connected component - the graph will reach consensus.

3 Sheaves on a graph and the sheaf Laplacian

Often when studying networks, we are interested not only in a simple state represented by a single number but a complex state represented by a vector in $\mathbb{R}^n$. For example in flocking, the state of each bird will be its velocity, or both position and velocity. In the field of algebraic topology, a sheaf describes data attached to a topological space. While this concept has historically been used only in purely theoretical mathematics, it can be borrowed in this context to attach data to a graph in a systematic way. We can think of a sheaf as a mathematical structure that organizes higher dimensional state vectors over nodes and connections of a graph.

A real, finite-dimensional, simplicial sheaf, which we will in general refer to as a sheaf, can be thought of as the space of "data" over the graph.[1] This means that to each vertex in V and each edge in E we assign a real, finite dimensional vector space.

$$\mathcal{F}(v_i) = \mathbb{R}^{n_{v_i}} \qquad \mathcal{F}(e_j) = \mathbb{R}^{n_{e_j}}$$

We call these vector spaces "stalks" of v_i or e_i, and in this setting we can think of them as state spaces associated to each vertex and edge. For example, in [9], the stalks over vertices represent "opinion spaces" for people in a social network and the stalks over edges represent "discourse spaces" where socially connected individuals compare their opinions. Whereas the state vector of a graph had one dimension per node, now we will have a stacked vector:

$$\vec{x} = \begin{bmatrix} \vec{x}_1 \\ \vec{x}_2 \\ \vdots \\ \vec{x}_{|V|} \end{bmatrix}$$

where $\vec{x}_i \in \mathbb{R}^{n_{v_i}}$. So this means $\vec{x} \in \mathbb{R}^{\sum n_{v_i}}$.

In addition to these spaces, whenever v_i is an vertex of e_j we assign *to that relationship* a linear map between their respective vector spaces. We denote this map as

$$\mathcal{F}(v_i < e_j) : \mathbb{R}^{n_{v_i}} \to \mathbb{R}^{n_{e_j}}.$$

A linear map between finite vector spaces can be represented as a matrix, so we also can write, $\mathcal{F}(v_i < e_j) = M_{i,j} \in \mathbb{R}^{n_{e_j} \times n_{v_i}}$. These maps are used to define what it means for the state of two vertices to "agree" or be compatible. If an edge e_k connects two vertices v_i and v_j, then the states of the vectors (choices of vectors in each vertex stalk) are said to agree if they map to the same vector in the edge stalk. If v_i is not a vertex of e_j, then we require $M_{i,j}$ to be the zero matrix.

We can then define a matrix that is analogous to the boundary matrix, which we call the sheaf boundary matrix. It is defined in block form:

$$D_{\mathcal{F}} = \begin{bmatrix} \pm M_{0,0} & \pm M_{0,1} & \dots & \pm M_{0,|V|} \\ \pm M_{1,0} & \pm M_{1,1} & \dots & \pm M_{1,|V|} \\ \vdots & \vdots & & \vdots \\ \pm M_{|E|,0} & \pm M_{|E|,1} & \dots & \pm M_{|E|,|V|} \end{bmatrix}$$

where the $\pm$ is chosen in the same way as the boundary matrix, so block row k will have all zero matrices, except if $e_k = (v_i, v_j)$ then $M_{i,k}$ is positive and $M_{j,k}$ is negative, depending on an arbitrary

[1] In the sense of category theory, it is a functor $\mathcal{F} : \textbf{Face}^{op}(G) \to \textbf{FinVect}_{\mathbb{R}}$.

orientation. Again analogously to the traditional boundary matrix, we define the **sheaf Laplacian** to be:

$$L_{\mathcal{F}} = D_{\mathcal{F}} D_{\mathcal{F}}^T$$

This matrix will not depend on the arbitrarily chosen orientation of the graph. It has been shown that the kernel of the sheaf Laplacian is a vector space of all possible "agreement states" for the sheaf. This space is also called the space of global sections. This lets us extend our consensus dynamics described before. While (2) modeled the flow of 1D opinions on a graph, the equation

$$\frac{d}{dt}\vec{x}(t) = -L_{\mathcal{F}}\vec{x} \tag{3}$$

will similarly model the flow to consensus on a sheaf. Any initial condition will flow to some steady state vector in the space of global sections, corresponding to a state where all vertices agree in the sense that they all map to the same vectors on edge stalks. In fact it can be shown that the steady-state vector is the projection of the initial condition onto the space of global sections.

4 Modeling flocking using sheaves on a graph

How do birds that flock together come to a consensus on which direction and speed to fly? If the birds only are aware of their immediate neighbors, how can coherent velocities be picked for the whole flock? Sheaves give us a language to model the problem. We model the position of each bird as a point in $\mathbb{R}^d$. For visualizations we use $\mathbb{R}^2$. We denote this point cloud as $B = \{\vec{b_i}\}_{i=0}^N$ where N is the number of birds in the flock. For now, the positions are constant. A graph and a sheaf over the graph are defined as a function of the birds' locations in space. Each bird is associated to a vertex b_i (note we use vector hats to denote the position in space, while this is the node of the purely combinatorial graph). We assume birds are only aware of other birds within a distance ϵ of themselves. This is modeled as an edge connecting these vertices.[2] We denote the resulting graph as G_B which has vertex set $V = \{b_i\}_{i=0}^N$ corresponding to each bird, and edge set $E = \{(b_i, b_j) \text{ s.t. } \|\vec{b_i} - \vec{b_j}\| < \epsilon \; \forall \; i < j < N\}$. The edges can be oriented arbitrarily.

We model the interactions of the birds as a sheaf. We denote the sheaf as $\mathcal{F}_B$ since it is once again a function of the point cloud B. Each bird has an "internal $\mathbb{R}^3$", modeled as the stalk over the birds respective vertex. Thus the sheaf $\mathcal{F}_B$ assigns to each vertex b_i a copy of $\mathcal{F}_B(b_i) = \mathbb{R}^3$. Think of this as the birds' internal compass or coordinate system. A choice of vector in the stalk of a vertex corresponds to that bird choosing a flight velocity, so we will refer to such vectors as $\vec{v_i}$. Each edge is then assigned $\mathcal{F}_B((b_i, b_j)) = \mathbb{R}$. This copy of $\mathbb{R}$ can be thought of as a way for the birds to compute their bearing in relation to the adjacent birds. Next, the sheaf assigns the linear functions $f_{i,j}(\vec{v}) = \mathcal{F}_B(b_i < (b_i, b_j)) : \mathbb{R}^3 \to \mathbb{R}$. This function is defined to be

$$f_{i,j}(\vec{v}) = \vec{v} \cdot \left(\frac{\vec{b_i} - \vec{b_j}}{\|\vec{b_i} - \vec{b_j}\|} \right)$$

for $i < j$ and otherwise $f_{j,i} = -f_{i,j}$. In words, this map gives the length of the projection of the chosen flight velocity of a bird onto the line between the bird and its neighbor. The flipping of the sign for maps $j < i$ is a convention so that when two birds connected by an edge are flying with the same velocity, the value of the push-forward of each vector to the edge stalk agrees.

A consequence of defining the sheaf this way is that if each bird picks the exact same flight velocity, this choice is in the space of global sections. Thus we say if the birds want to fly "coherently" they will choose velocities in the space of global sections. In general however, there are interesting "coherent velocities" that are not the flock moving exactly the same direction.

[2]This is the 1-skeleton of the Cech complex at parameter $\epsilon/2$

5 Consensus dynamics of a flock

Given any vector $\vec{v}$ in $\mathbb{R}^{|V|*d}$ where d is the dimension of the flock, we can think of this as a "random" choice of velocity for each bird. This will not necessarily be coherent. However, if we consider this vector as the initial condition for the system (3), the limiting vector will give the "closest" coherent direction for the flight paths. Figure 1 shows an initial random vector, and the final vector to which it flows.

This flow shows that using only information about their immediate neighbors, each bird is able to choose a direction that is coherent in the sheaf sense. This is an extension of the graph consensus model to account for a more nuanced definition of what "agreement" means. The steady state vectors of this system all lie in the kernel of the Laplacian matrix - the space of global sections.

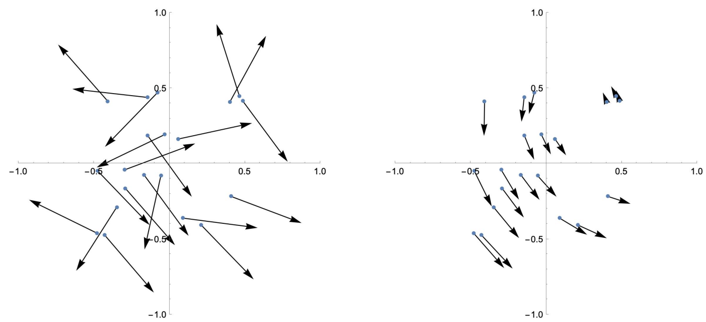

Figure 1: On the left, a randomly chosen element of the state space. On the right, the projection of this vector onto the space of global sections. This is the where the initial condition on the left converges to while undergoing the consensus dynamics.

6 The global sections of a flock sheaf

We explore the global sections of the flock sheaf - what choices of velocities can the birds pick such that they are in agreement? We focus mostly on 2D flocks for ease of visualization. Figures 2, 3, and 4 show examples of 2D flock graphs, followed by visualizations of vectors which span the spaces of global sections.

In figure 2, the graph is well connected. The space of global sections in this case is 3-dimensional. The first two vectors show that any choice of state vectors where the birds move exactly in the same direction is in the space of global sections. This is as expected. However, the third dimension shows that any choice of state vector where the birds pick directions corresponding to rotations about their center of mass is also a coherent direction. Since this is a vector space, any linear combination of these three vectors is also a coherent direction. These look like rotations about any arbitrary point in the plane.

In figure 3, the graph is disconnected, with two components. The space of global sections in this case is 6-dimensional. Four of these vectors show that any state vector where all the birds in a component agree on exactly the same direction, but not necessarily the same direction in each component, is a coherent choice. Then each component can rotate independently as well, as in the single component before.

In figure 4, the graph is connected, but with two main components that are connected via a single vertex. The space of global sections in this case is 4-dimensional. As before the whole flock can move or rotate together, accounting for 3 dimensions. However, the fourth dimension shows that the two components are able to rotate independently about the vertex connecting them.

7 Eigenvectors of the flock sheaf Laplacian

The graph Laplacian is used commonly in spectral graph theory to study the properties of graphs [3]. While the kernel tells us about the connected components, the eigenvectors corresponding to non-zero

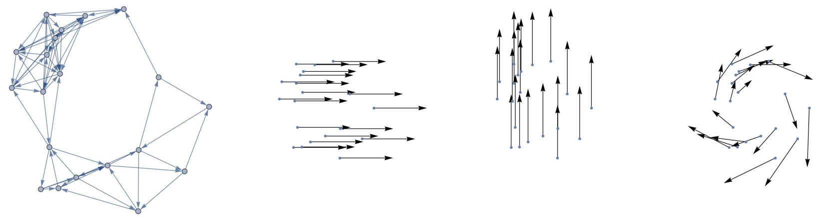

Figure 2: The first panel shows the underlying graph (indicating which birds are "close enough" to be aware of each other). The following three panels visualize vectors that span the space of global sections for this flock sheaf. The space of global sections is 3-dimensional in this case.

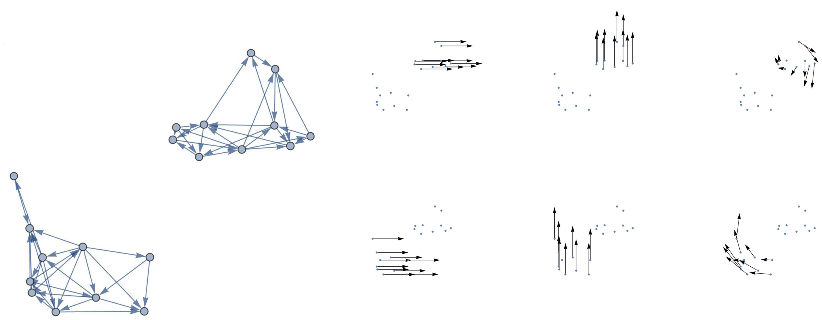

Figure 3: The first panel shows the underlying graph. The following 6 panels visualize vectors that span the space of global sections for this flock sheaf. The space of global sections is 6-dimensional in this case. Notice 3 dimensions each correspond to a connected component.

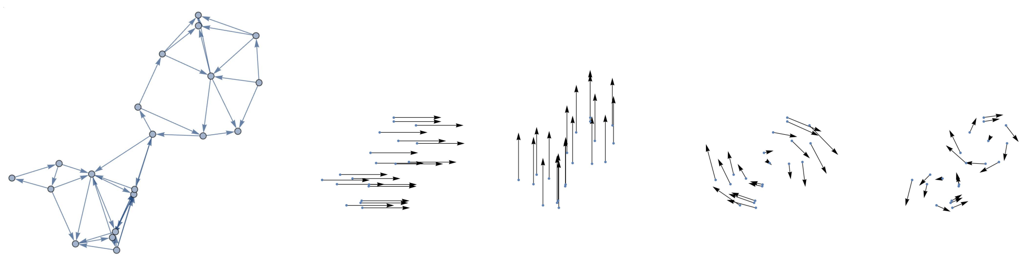

Figure 4: The first panel shows the underlying graph. The following 4 panels visualize vectors that span the space of global sections for this flock sheaf. The space of global sections is 4-dimensional in this case. Notice the last dimension corresponds to ways the two components are free to rotate independently of each other.

eigenvectors give connectivity information about the graph as well. Most famously, the eigenvector corresponding to the first non-zero eigenvalue is called the Fiedler vector, and it is known to give information about the smallest cut necessary to split a graph into disconnected components [4].

Figure 5 visualizes the eigenvectors of a connected flock sheaf. The first non-zero eigenvalue seems to contain similar information to the Fiedler vector, as the magnitude of the vector on the least-connected point seems much larger than the rest of the vectors. Figure 6 shows the eigenvectors of a flock of points lying on a circle. This emphasizes a connection between the discrete Laplacian and the continuous Laplacian, in that the eigenvectors are reminiscent of Fourier modes. Future work should explore this connection further.

8 Flocking dynamics

The consensus dynamics of a flock show that for a given configuration of birds, the birds can find a coherent direction to fly by flowing along the ODE defined by the sheaf Laplacian. This does not

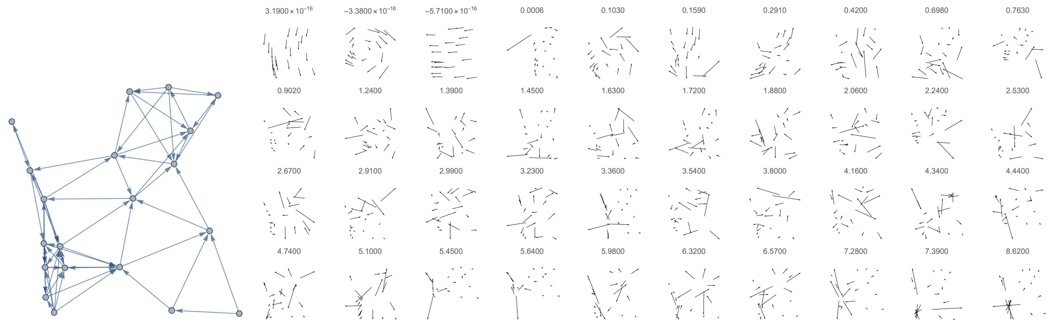

Figure 5: For the 2D flock on the left, we visualize the eigenvectors of the sheaf Laplacian. The corresponding eigenvalues are shown above each visualization. Notice the first three eigenvalues are numerically zero, meaning these three vectors give a basis for the kernel, i.e. the space of global sections. The first non-zero eigenvalue seems to contain information similar to the Fiedler vector of a graph Laplacian.

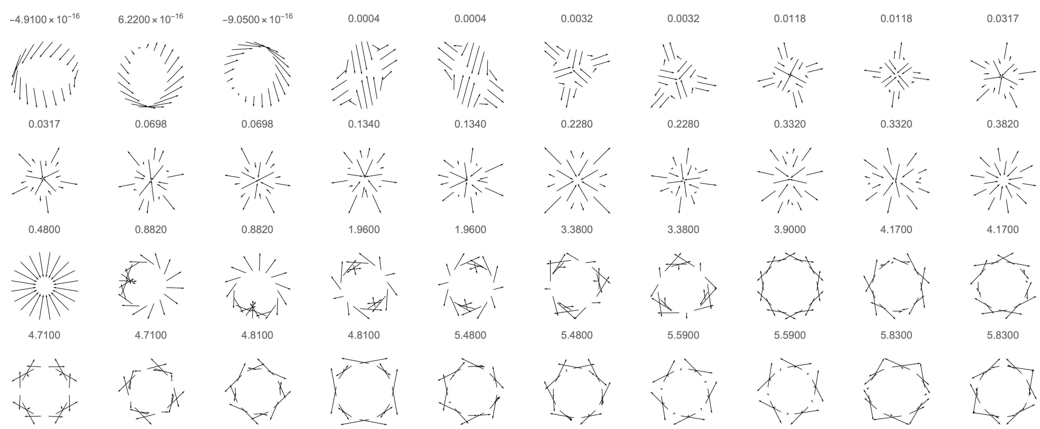

Figure 6: If a 2D flock is given as 20 points evenly distributed on the unit circle with each vertex connected two its 4 nearest neighbors, these visualize the eigenvectors of the sheaf Laplacian. The corresponding eigenvalues are shown above each visualization.

take into account that once the birds move the graph and sheaf itself will change. To account for this we can define dynamics on the positions of the birds, which will be affected by the sheaf dynamics. The sheaf will then be affected by the changing position of the birds. We define the dynamics to be:

$$\frac{d}{dt}\vec{b} = \vec{v} \tag{4}$$

$$\frac{d}{dt}\vec{v} = -\alpha L_{\mathcal{F}}(\vec{b})\vec{v} \tag{5}$$

where $\vec{b}$ is the stacked vector of all the birds positions. Note that $L_{\mathcal{F}}(\vec{b})$ depends non-linearly on the birds locations. α is a parameter that controls the "reaction time" of the birds, by emphasizing the velocity correction. This is a dynamical system on an adaptive network - the network affects the dynamics and the dynamics affect the network in turn. Figures 7 and 8 show numerical solutions of this system, computed on a standard personal computer.

This is a dynamical system with (number of birds) $*d*2$ variables, where d is the dimension of the flock. Each bird has d position state variables and d velocity state variables. This dynamical system cannot be framed quite as simply using sheaves as the consensus sytem (3). However, the fact that at all times, the matrix $L_{\mathcal{F}}(\vec{b})$ is a sheaf Laplacian could prove useful for further analysis.

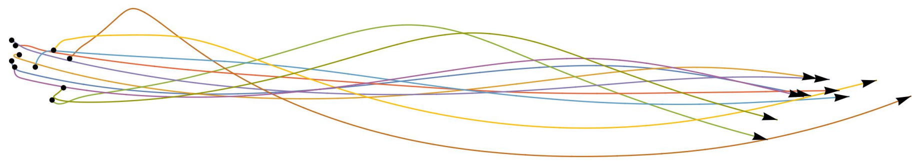

Figure 7: An example of a flock evolving using the dynamics in (4) and (5). The initial positions were chosen randomly to be within the unit square, and the velocities were picked randomly to have entries between -1/2 and 1/2. This was simulated with using a forward Euler method. Parameters were chosen $\alpha = 0.5$, $\epsilon = 0.8$, $\Delta t = 0.05$, and 2400 timesteps were taken.

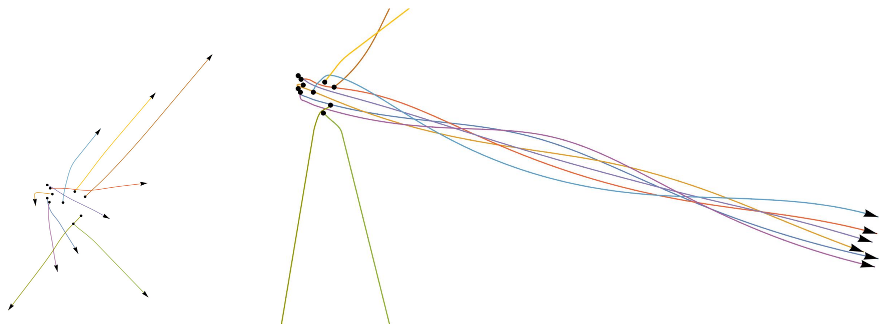

Figure 8: These flocks have the same initial positions and velocities as in figure 7. However the flock shown on the left α was changed to 0.1, and only 200 timesteps are shown. For the flock on the right, ϵ was changed to 0.6. This shows that the initial velocities can push birds away from the others before the Laplacian dynamics encourage consensus velocities, and some birds will not join the "flock". Once birds are more than ϵ away from all other birds, they will move straight indefinitely.

9 Conclusion and discussion

We defined a model of animal flocking using the theory of sheaves over a graph. The theory of sheaves proved useful to create a model of birds finding coherent flight paths via analogy to opinion consensus on a social network. Then true flocking was modeled by letting the positions and velocities evolve simultaneously.

This model is an example of how sheaves organize data over networks. By defining maps between the stalks above vertices and edges locally, the global sections of the sheaf describe the space of states that satisfy all the local maps. Our example of flocking shows how this structure is helpful in the modeling of emergent phenomena where global behavior arises from local rules.

To our knowledge, flocking models using the sheaf Laplacian in the way we define here have not been previously investigated. In future work more quantitative comparisons of this model and other existing flocking models should be explored. Our dynamical system is very similar to those described in [15, 17], but neglecting terms that model attraction and repulsion of the agents. However, these models use a matrix defined as the Kronecker product of the graph Laplacian and a $d \times d$ identity matrix in place of the sheaf Laplacian. The kernel of the matrix used in these other models consists of choices of velocity where every bird agrees exactly on a velocity for each connected component. The fact that the kernel of the sheaf Laplacian has additional degrees of freedom gives our system qualitatively different behavior from those other systems. In the preprint [6], similar consensus dynamics are explored using cellular sheaves and sheaf Laplacians in the context of coordinating multi-agent systems. However, in this case the network topology is fixed and does not depend on each agents' position. Further comparison of our model and these other works, as well as comparison to empirical data, could show strengths and weaknesses of the proposed model.

Additional future work could explore more dynamical systems on sheaves over networks. Applying the machinery of sheaves to other models could lead to interesting analysis. Higher order networks modeled as simplicial complexes have Hodge Laplacians which contain information about higher-dimensional connectivity. Another direction of future work would be to explore ways of assigning sheaves to such structures and exploring the sheaf Hodge Laplacians.

Acknowledgments and Disclosure of Funding

Thanks to the CSU Topology Seminar for inspiring this work. In particular thanks to Dr. Chris Peterson, Dr. Amit Patel, Dr. Margaret Cheney, and Dr. Clay Shonkwiler for many discussions and revisions. Thanks as well to Michael Moy, Tatum Rask, and Jacob Cleveland for always being willing to talk topology.

References

[1] Rico Berner, Thilo Gross, Christian Kuehn, Jürgen Kurths, and Serhiy Yanchuk. Adaptive dynamical networks. *Physics Reports*, 1031:1–59, August 2023.

[2] B. Cessac. A view of neural networks as dynamical systems. *International Journal of Bifurcation and Chaos*, 20(06):1585–1629, 2010.

[3] F.R.K. Chung. *Spectral Graph Theory*. Conference Board of Mathematical Sciences. American Mathematical Society, 1997.

[4] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.

[5] Robert W Ghrist. *Elementary applied topology*, volume 1. Createspace Seattle, 2014.

[6] Tyler Hanks, Hans Riess, Samuel Cohen, Trevor Gross, Matthew Hale, and James Fairbanks. Distributed multi-agent coordination over cellular sheaves, 2025.

[7] Jakob Hansen. *Laplacians of Cellular Sheaves*. PhD thesis, University of Pennsylvania, 2020.

[8] Jakob Hansen and Robert Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3(4):315–358, December 2019.

[9] Jakob Hansen and Robert Ghrist. Opinion dynamics on discourse sheaves. *SIAM Journal on Applied Mathematics*, 81(5):2033–2060, 2021.

[10] A. Hatcher. *Algebraic Topology*. Algebraic Topology. Cambridge University Press, 2002.

[11] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.

[12] A. Jadbabaie, Jie Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.

[13] Y. Kuramoto. *Chemical Oscillations, Waves, and Turbulence*. Dover books on chemistry. Dover Publications, 2003.

[14] Leonie Neuhäuser, Andrew Mellor, and Renaud Lambiotte. Multibody interactions and nonlinear consensus dynamics on networked systems. *Physical Review E*, 101:032310, Mar 2020.

[15] R. Olfati-Saber. Flocking for multi-agent dynamic systems: algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.

[16] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '87, pages 25–34, New York, NY, USA, August 1987. Association for Computing Machinery.

[17] H.G. Tanner, A. Jadbabaie, and G.J. Pappas. Stable flocking of mobile agents part i: dynamic topology. In *42nd IEEE International Conference on Decision and Control (IEEE Cat. No.03CH37475)*, volume 2, pages 2016–2021 Vol.2, 2003.

A Graph theory background

Definition: A **graph** G is a set of vertices $V = \{v_1, v_2, ..., v_N\}$ and a set of edges E indicating connections between any two vertices, i.e. $E \subset \{\{v_i, v_j\} | v_i, v_j \in V\}$. We say $G = (V, E)$

The vertices would represent the nodes of the network - individual people in a social network, producers/consumers in a power network, or species in an ecological network. The edges represent the connections - friendships in a social network, or transmission lines in a power network, or predator/prey relationships in an ecological network. The graph lets us study the network through a mathematical framework.

Graphs come in different kinds. *Directed graphs* add directionality to the edges, indicating the "head" and "tail" of the edge. The only change to our definition would be that our edges are tuples (v_i, v_j) where order matters, unlike with sets. These kinds of graphs can be useful in modeling networks where directionality is important, for example a plumbing network, as water can only flow one way. *Weighted graphs* assign to each edge a "weight" which we will denote w_{ij}, indicating it is assigned to the edge $\{v_i, v_j\}$. If that edge is not in the graph, w_{ij} is assumed to be zero. Perhaps this could model the transmission capability of the wires connecting power stations. These kinds of graphs add more complexity to the model, but in turn can model more complex networks. The key in mathematical modeling is balancing the complexity of the model with the required accuracy needed.

An important concept in graph/network theory is the "degree" of a node.

Definition the **degree** of a node v_i, denoted $\deg(v_i)$ is the number of edges adjacent to the vertex. For weighted graphs $\deg(v_i) = \sum_{j=1}^{N} w_{ij}$

While the set-notation definition of a graph is enough to represent any graph, often it is useful to represent a graph in the form of an adjacency matrix. Matrices are easily represented in a computer, so adjacency matrices are useful for computation. Frequently they are sparse, and sparse matrix algorithms can greatly speed up such computations.

Definition: The **adjacency matrix** A for a graph G is a matrix with elements

$$a_{ij} = \begin{cases} 1 & \text{if } \{v_i, v_j\} \in E \\ 0 & \text{otherwise} \end{cases}$$

This definition can be adjusted to represent either a directed or weighted graph as well.

Definition: The **weighted adjacency matrix** A for a weighted graph G with edge weights w_{ij} is a matrix with elements

$$a_{ij} = \begin{cases} w_{ij} & \text{if } \{v_i, v_j\} \in E \\ 0 & \text{otherwise} \end{cases}$$

B Higher order networks and simplicial homology

In reality, not all interactions and connections involve only two parties (we call two way interactions/relationships *dyadic*). Perhaps this is most obvious in a social network. While the knowledge that "person A is friends with person B" is important to understanding a social dynamic, knowing that "person A, person B, and person C are all in the same family" or "persons A, B, C, and D all attend the same school" give you even more detailed information. If real-world networks include such *supra-dyadic* or *higher-order* relationships, the mathematician needs tools to model such phenomena. This gives rise to two related ways to abstractly model a higher-order network: hypergraphs and simplicial complexes.

A hypergraph is the most flexible extension of a graph. It has no restrictions on the order or types of interactions that can be modeled. It simply tells you a list of vertices and any interactions between them.

Definition: A **hypergraph** H is a set of vertices $V = \{v_1, v_2, ..., v_N\}$ and a set of edges E indicating connections between vertices, i.e. $E \subset \mathcal{P}(V)$ where $\mathcal{P}(V)$ is the power set of V. We say $H = (V, E)$

One could split up the set of "edges" by the order of the interaction, for example $E_1 \subset \{\{v_i, v_j\} | v_i, v_j \in V\}$ are the usual dyadic edges, $E_2 \subset \{\{v_i, v_j, v_k\} | v_i, v_j, v_k \in V\}$ represents 3-way interactions and so on. Then $E = \cup_{i=1}^{N} E_i$

A simplicial complex is traditionally used more often in the field of topology. It is often thought of as a way to break up or "triangulate" a topological space in a way that allows us to use combinatorial and algebraic arguments to uncover properties of a space. However, as a model of higher-order networks, we can reframe the definition of a simplicial complex as a hypergraph with added structure on the collection of edges. While any arbitrary connections are permitted in a hypergraph, there are additional rules on our edge-set for a hypergraph to be considered a simplicial complex.

Definition: A **simplicial complex** $K = (V, E)$ is a hypergraph whose edge set is closed under inclusion.

This means E has the property that if an edge e is in E then any subset of e must also be in E. Using standard notation from algebraic topology, we will refer to edges as σ rather than e. If σ_i is a face of σ_j, we indicate this relation as $\sigma_i < \sigma_j$.

The field of Algebraic Topology studies how we can assign algebraic structures to topological spaces in order to study them and tell them apart. As simplicial complexes are one of the most understood types of topological spaces, we can use knowledge from this field to understand more about higher-order networks when simplicial complexes are used as an appropriate model.

We focus here exclusively on simplicial homology with real coefficients. For a full treatment on homology with general coefficients, or other homology theories like singluar homology see [10].

We call the hyper-edges "k-simplices" where $k + 1$ is the number of vertices in the edge. A 0-simplex is simply a vertex. The subsets of a k-simplex are called its faces. Faces that are not the k-simplex itself are proper faces. It is standard to use σ rather than e to represent a simplex. Let X be a finite, oriented, n-dimensional, simplicial complex with m simplices in total. Let m_d be the number of simplices of dimension d so that $m = m_0 + m_1 + \cdots + m_n$. We introduce two notations for indexing simplices in X. We denote a generic simplex in X as σ and write

$$X = \{\sigma_i\}_{i=0}^{(m-1)}.$$

This notation is used when we want to talk about all simplices and do not care about the dimensions of the individual simplex.

If a simplex is of dimension d we write σ^d, or when a more concrete simplex is specified we write $[v_0 \dots v_d]$. We denote the set of d-dimensional simplices as

$$E_d = \{\sigma_i^d\}_{i=0}^{(m_d-1)}$$

so that

$$X = \bigcup_{d=0}^{n} E_d$$

It is often useful to keep track of the "orientation" of a k-simplex. We define

Definition An **oriented k-simplex** is a tuple of $k + 1$ vertices, written $\sigma = [v_0, ..., v_k]$.

The **orientation** of an oriented simplex is either positive or negative. The orientation $+$ or $-$ is in general not meaningful in itself, but must be kept track of to indicate if two simplices consisting of the same vertices are oriented the same way or oppositely. Two simplices consisting of the same vertices have the same orientation if they are an even permutation of each other, and opposite orientations if they are odd permutations of each other. For example we write $[v_0, v_1] = -[v_1, v_0]$. An **oriented**

simplicial complex is a simplicial complex where each simplex has been assigned an orientation. This can be done canonically by numbering the vertices, and defining the simplex with vertices listed from lowest to highest index (or even permutations of this ordering) as the positive orientation.

There are various ways to represent a simplicial complex or hyper-graph in a data-structure like adjacency matrices. For example, a simplicial complex can be represented as an "adjacency matrix" for each dimension, indicating how k-simplices are connected to their $k-1$ simplex faces. This idea is similar to the boundary matrices explained below.

Definition: The k^{th} **chain group** with real coefficients of an oriented simplicial complex X, or simply the k^{th} chain group of X, is the vector space $\mathcal{C}_k(X)$ representing formal sums of oriented simplices with real coefficients.

Definition The **boundary operator** $\partial_k : \mathcal{C}_k(X) \to \mathcal{C}_{k-1}(X)$ is the linear map on elements of the kth chain group given by its action on simplices:

$$\partial([v_0, ..., v_k]) = \sum_{i=1}^{k} (-1)^k \, [v_0, ..., \hat{v}_i, ..., v_k]$$

By definition $\partial([v]) = 0$ for any 0-simplex. The boundary operator has the property that $\partial(\partial(\sigma)) = 0$.

Definition The k^{th} **boundary group** $\mathcal{B}_k(X)$ is the image of the boundary operator ∂_{k+1}

Definition The k^{th} **cycle group** $\mathcal{Z}_k(X)$ is the kernel of the boundary operator ∂_k

Definition: The k^{th} **homology group** is $\mathcal{H}_k(X) = \mathcal{Z}_k(X)/\mathcal{B}_k(X)$

Elements of the boundary group are called boundaries. Elements of the cycle group are called cycles. Notice that because of the property $\partial(\partial(\sigma)) = 0$, we know $\mathcal{B}_k \subseteq \mathcal{Z}_k \subseteq \mathcal{C}_k$. Elements of the homology group are equivalence classes of of k-cycles that do not bound any $k+1$ chains.

The chain group $\mathcal{C}_k(X)$ is isomorphic to $\mathbb{R}^{|E_k|}$ so we can think of elements of this space either as formal sums of simplices, or as column vectors. In this framing, the linear boundary operator can be thought of a matrix. The **boundary matrix** $B_k \in \mathbb{R}^{|E_{k-1}| \times |E_k|}$ is the matrix representation of the boundary operator. The columns of the matrix are the vectors in $\mathbb{R}^{|E_{k-1}|}$ that correspond to the boundary operator applied to each simplex in E_k

Definition The n$^{\text{th}}$ **Hodge Laplacian** is the matrix

$$L_n = W_{n+1} B_n^T W_n^{-1} B_n + B_{n+1} W_{n+2} B_{n+1}^T W_{n+1}^{-1}$$

Where B_n are matrix representations of the boundary operators and W_n are diagonal weight matrices. In the case that $n = 0$ and $W_{n+2} = W_{n+1} = I$, we have $L_0 = B_1 B_1^T$ which can be shown to exactly recover the graph Laplacian. This extends the concept of a Laplacian to simplicial complexes. It has been shown that the chain group can be decomposed as

$$\mathcal{C}_k(X) \cong \text{im}(B_k^T) \oplus \text{im}(B_{k+1}) \oplus \ker(L_k)$$

and

$$\ker(L_k) \cong \ker(B_k)/\text{im}(B_{k+1}) = \mathcal{H}_k(X)$$

So the Hodge Laplacians tell us about our homology. In particular, the dimension of the kernel of the k^{th} Laplacian is the betti number β_k.

B.1 Dynamical Systems on Higher Order Networks

As discussed, not all types of networks can be modeled with a graph. "Higher-Order Networks" include connections between nodes that are not just dyadic. What would be the appropriate way to model dynamics on such a network? A three-way connection may contribute differently to the system than a two-way connection. Therefore we would model generally the dynamics of the state of each vertex as

$$\dot{x}_k = F(x_k) + \sum_{j=1}^{N} a_{jk} G(x_k, x_j) + \sum_{j=1}^{N} \sum_{l=1}^{N} a_{jkl} H(x_k, x_j, x_l) + ...$$

This is one way to think of the networks as "higher order", the dynamics on the graph is just a truncation of this series, similar to a truncated taylor series.

The language of homology lets us think of our state vector as an element of the 0th chain group $\mathcal{C}_0(X)$, as each entry of the vector corresponds to a number on each vertex. However, we are not always interested in the state on a vertex, but rather sometimes we are concerned with the state on an edge, or on a higher order connection. This is useful when modeling network flows, or functions in 3D volumes, and similar phenomena. In this case we can simply think of our state vector as living in $\mathcal{C}_k(X)$, and define similar systems.

B.1.1 Consensus on Higher Order Networks

The dynamical system on a simplicial complex X, modeling opinion dynamics as before can now be written:

$$\frac{d}{dt}\left[\vec{x}\right] = L_0 \vec{x}$$

On a theoretical level, we think of the dynamics as happening in the 0th chain group of the graph (a 1-dimensional simplicial complex), i.e. $\vec{x} \in \mathcal{C}_0(G)$. If we have higher order connections, i.e. $\vec{x} \in \mathcal{C}_0(X)$ where X is a simplicial complex rather than a graph, the dynamics will not change, since L_0 will never depend on any higher order connections, since the 0th Laplacian is dependent only on the 1-skeleton of the complex. In fact, it has been shown that if any linear higher-order interactions are introduced, the system can simply be rewritten as a graph with different weights. [14]. Non-linearities are necessary to have interesting differences between consensus on a graph and consensus on a higher order network.

However, sometimes we are interested not in consensus on the nodes of a higher order network but instead on the edges, or triangles, or higher-order connections/simplicies. If we think of $\vec{x} \in \mathcal{C}_k(X)$ for some simplicial complex X, we can model "consensus" in the chain group by the differential equation

$$\frac{d}{dt}\left[\vec{x}\right] = -L_k \vec{x}$$

We used many properties of the graph Laplacian to infer what the solutions of this dynamical system will look like in the 0th case. Do the same properties always hold for Hodge Laplacians? This is not always the case, however the knowledge of Hodge decomposition tell us that the equilibrium solutions will be in a subspace of the chain group of dimension equal to the betti number of the simplicial complex. We observe the differences between the following two examples.

B.2 Sheaf Cohomology

The co-homology theory relevant to this discussion is a special case of cellular homology. Rather than defining cellular sheaves in their full generality we restrict our attention to the relevant case which, in addition to being more straight-forward, is also easily computable using standard methods in linear algebra.

A real, finite-dimensional, simplicial sheaf, which we will in general refer to as a sheaf, can be thought of as the space of "data" over the simplicial complex. In the sense of category theory, it is a functor $\mathcal{F} : \mathbf{Face}^{op}(X) \to \mathbf{FinVect}_{\mathbb{R}}$. This means that to each simplex in X we assign a real, finite dimensional vector space.

$$\mathcal{F}(\sigma_i) = \mathbb{R}^{n_i}$$

We call this vector space the "stalk" of σ_i. In addition, whenever σ_i is a face of σ_j we assign to that relationship a linear map between their respective vector spaces. We denote this map as

$$\mathcal{F}(\sigma_i < \sigma_j) : \mathbb{R}^{n_i} \to \mathbb{R}^{n_j}.$$

We also require that all maps commute, i.e. that if $\sigma_i < \sigma_j < \sigma_k$ then the following holds:

$$\mathcal{F}(\sigma_i < \sigma_k) = \mathcal{F}(\sigma_j < \sigma_k) \circ \mathcal{F}(\sigma_i < \sigma_j)$$

Note that if we let simplices be faces of themselves, then this condition means $\mathcal{F}(\sigma_i < \sigma_i) = \text{Identity}$. Because of the conditions of linearity and finiteness, we can describe $\mathcal{F}(\sigma_i < \sigma_j)$ as an $n_i \times n_j$ matrix. With the additional constraint that maps commute, all the information we need is contained in the matrices describing maps from d-dimensional simplices to $(d+1)$-dimensional simplices. We denote the matrix describing $\mathcal{F}(\sigma_i^d < \sigma_j^{d+1})$ as $M_{i,j} \in \mathbb{R}^{n_j \times n_i}$. Note that using this notation, we have $i \in \{1, 2, ..., m_d\}$ and $j \in \{1, 2, \ldots, m_{d+1}\}$. As a convention, if σ_i^d and σ_j^{d+1} are not adjacent, then we define $M_{i,j}$ to be the zero matrix.

We define the d-dimensional cochain group as the direct sum of the vector spaces corresponding to the d-dimensional simplices of X.

$$C^d(X, \mathcal{F}) = \bigoplus_{i=0}^{m_d - 1} \mathcal{F}(\sigma_i^d) \simeq \mathbb{R}^{N_d}$$

Where N_d is the sum of all n_i corresponding to σ_i^d. We think of elements of the cochain groups as a choice of element for each stalk $\mathcal{F}(\sigma_i^d)$. pick a point $\vec{v}_i$ in each $\mathbb{R}^{n_i} = \mathcal{F}(\sigma_i^0)$ we represent the corresponding vector $\vec{v} \in \mathbb{R}^{N_0}$ as the stacked vector

$$\vec{v} = \begin{bmatrix} \vec{v}_0 \\ \vec{v}_1 \\ \vdots \\ \vec{v}_{m_0} \end{bmatrix}$$

We then have a cochain complex

$$0 \xrightarrow{D_{-1}} C^0(X, \mathcal{F}) \xrightarrow{D_0} C^1(X, \mathcal{F}) \xrightarrow{D_1} \ldots \xrightarrow{D_{n-2}} C^{n-1}(X, \mathcal{F}) \xrightarrow{D_{n-1}} C^n(X, \mathcal{F}) \xrightarrow{D_n} 0$$

Where each map D_k is defined as block matrices similar to $D_{\mathcal{F}}$ in the main paper. In this case $D_{\mathcal{F}} = D_0$

Where the signs are determined by the orientation of the simplicial complex. Note that $D_k D_{k+1} = 0$

The kth sheaf cohomology of X is defined to be

$$H^k(X, \mathcal{F}) = \text{Ker}(D_k)/\text{Im}(D_{k-1})$$

A choice of vectors $\vec{v}_{\sigma_i}$ in $\mathcal{F}(\sigma_i)$ for each simplex in the complex is called a "global section" of $\mathcal{F}$ if it satisfies the sheaf in the appropriate way: if $\vec{v_{\sigma_j}} = \mathcal{F}(\sigma_i < \sigma_j)(\vec{v_{\sigma_i}})$ for all $\sigma_i < \sigma_j$. The set of all global sections of a sheaf can be shown to be a vector space itself. We denote it $\Gamma(X, \mathcal{F})$. It can also be shown that $H^0(X, \mathcal{F}) \simeq \Gamma(X, \mathcal{F})$. It may seem surprising that the global sections depend only on $C^0(X, \mathcal{F})$, $C^1(X, \mathcal{F})$, and D_0, but this is a consequence of the requirement that maps commute within the sheaf. [5]

We define the kth up-Laplacian and down-Laplacian respectively as

$$L_k^{\text{up}} = D_k^T D_k$$

$$L_k^{\text{down}} = D_{k-1} D_{k-1}^T$$

We then define simply the kth Laplacian as

$$L_k = L_k^{\text{up}} + L_k^{\text{down}}$$

It has been proven that analogously to the standard Hodge decomposition we have

$$\text{Ker}(L_k) \simeq H^k(X, \mathcal{F})$$

and in particular

$$\text{Ker}(L_0) \simeq H^0(X, \mathcal{F}) \simeq \Gamma(X, \mathcal{F})$$

This fact eliminates the need to compute the more difficult quotient space $H^k(X, \mathcal{F})$ by reducing the problem to finding the null space of a single matrix. It is convenient that the kernel of the Laplacian is a concrete subspace of $C^k(X, \mathcal{F})$. This gives us canonical or "harmonic" cycles C^k which correspond directly to equivalence classes in H^k

C Derivation of consensus on a network

We show a derivation of the model for opinion consensus on a network. We consider a graph G which models a social network. Each individual, out of N individuals total, is represented by a node v_i, and the edges are weighted with positive numbers w_{ij} representing the strength of the bond/friendship between individuals v_i and v_j. For two unconnected individuals, w_{ij} is zero, so we have a weighted adjacency matrix $A = (w_{ij})$. Each individual has a current state/opinion, represented by a real number. We are interested in the evolution of this opinion as time progresses so we say individual v_i's opinion at time t is $x_i(t)$

It has been said that "you are the average of the 5 people you spend the most time with". We take this to heart and say, that an individual i will adjust their opinion to be closer to their neighbor j with rate w_{ij}. We can write this mathematically as

$$x_i(t + \Delta t) = x_i(t) + \sum_{j=1}^{N} \Delta t w_{ij}(x_j(t) - x_i(t))$$

Rearanging and taking the limit as Δt goes to zero, we find

$$\lim_{\Delta t \to 0} \frac{x_i(t + \Delta t) - x_i(t)}{\Delta t} = \frac{d}{dt}[x_i(t)] = \sum_{j=1}^{N} w_{ij}(x_j(t) - x_i(t))$$

If we rearrange this expression we see

$$\frac{d}{dt}\left[x_i(t)\right] = -x_i \sum_{j=1}^{N} w_{ij} + \sum_{j=1}^{N} w_{ij} x_j$$
$$= -(x_i \deg(v_i) - \sum_{j=1}^{N} w_{ij} x_j)$$

We notice that if we write out our entire state vector $\vec{x} = (x_i)$ we have

$$\begin{aligned}\frac{d}{dt}\left[\vec{x}(t)\right] &= -(D\vec{x} - A\vec{x}) \\ &= -(D - A)\vec{x} \\ &= -L\vec{x}\end{aligned}$$

D Examples

Example 1:

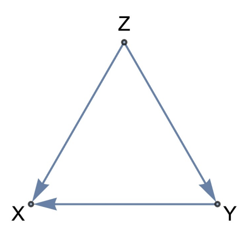

We consider the 1-D simplicial complex (graph) on three vertices. Orientation is indicated in the figure. Both $\mathcal{C}_0(X)$ and $\mathcal{C}_1(X)$ are isomorphic to $\mathbb{R}^3$. We denote an element of $\mathcal{C}_0(X)$ as $\vec{x}$ and $\mathcal{C}_1(X)$ as $\vec{y}$ We only have one non-zero boundary matrix B_1. The boundary operator can be written $\partial_1(\vec{y}) = B_1\vec{y} \in \mathcal{C}_0(X)$

$$B_1 = \begin{bmatrix} 1 & 0 & 1 \\ -1 & 1 & 0 \\ 0 & -1 & -1 \end{bmatrix}$$

Then we have $L_0 = B_1 B_1^T$ and $L_1 = B_1^T B_1$ below

$$L_0 = \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix} \qquad L_1 = \begin{bmatrix} 2 & -1 & 1 \\ -1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}$$

We consider first, the normal consensus on vertices of the graph. This is exactly the dynamcial system described in the previous section.

$$\frac{d}{dt}\left[\vec{x}\right] = -L_0\vec{x}$$

Since L_0 is the graph Laplacian, and our graph is connected, we know that 0 is an eigenvector and that the kernel is span$\{[1,1,1]^T\}$. The general solution to this equation is

$$\vec{x}(t) = c_1 \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} + c_2 e^{-3t} \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} + c_3 e^{-3t} \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix}$$

for any choice of c_1, c_2, c_3 determined by an initial condition. You can see that as $t \to \infty$, this converges exponentially to a point on the vector $[1,1,1]^T$, where the state of every vertex "agrees" - the verticies are in consensus.

Now what about consensus on the edges? We consider the extremely similar system

$$\frac{d}{dt}[\vec{y}] = -L_1\vec{y}$$

the matrix L_1 is a Hodge Laplacian (not a graph Laplacian) we are not guaranteed the same facts we discussed previously. We find the solution to this system is

$$\vec{y}(t) = c_1 \begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix} + c_2 e^{-3t} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} + c_3 e^{-3t} \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix}$$

This is very similar to before, but now the kernel of the Laplacian is span$\{[-1-,1,1]^T\}$, and the limit of each solution is a point on this line. Do the negatives indicate that our edges "disagree"? No - the agreement takes into account the orientation, and a stable solution is when the states are a harmonic cycle - representing a point in $\mathcal{H}_1(X)$, or a rotation around the central hole. What happens if we fill in the hole with a triangle?

Example 2:

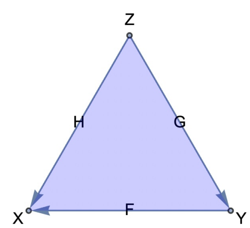

We now consider a very similar simplicial complex, but with the added 2-simplex in the center. The 1-skeleton of the complex has not changed, therefore L_0 will be identical in this situation to the previous example, as will the consensus dynamics. However, there is now a non-trivial boundary matrix:

$$B_2 = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}$$

And now we will have

$$L_1 = B_1^T B_1 + B_2 B_2^T = \begin{bmatrix} 2 & -1 & 1 \\ -1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} + \begin{bmatrix} 1 & 1 & -1 \\ 1 & 1 & -1 \\ -1 & -1 & 1 \end{bmatrix} = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{bmatrix}$$

This matrix does not have 0 for an eigenvalue - it has eigenvalue 3 repeated three times. The system

$$\frac{d}{dt}[\vec{y}] = -L_1\vec{y}$$

has general solution

$$\vec{y}(t) = e^{-3t} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}$$

So any solution converges to the stable node at $\vec{y} = 0$. Hodge decomposition tells us this to begin with - the simplicial complex has no β_1, so the dimension of the kernel of the Laplacian is zero. In some sense we can think of this in terms of flows on the edges. In the previous example, there are coherent ways for flows to cycle around the middle hole, while in this case if the hole is filled in the only coherent steady state flow is none at all.

TAG-DS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The introduction and abstract accurately summarize and reflect on the problem of modeling flocking and background. The contributions and scope of the paper are explained, mainly a reframing of the problem using the construction of sheaves.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper discusses simple dynamical system models of flocking, viewed through a sheaf theory lens. The authors do not claim that this model is extraordinarily novel or better than existing flocking models, but rather that the use of sheaf theory can provide insights that are more difficult to articulate without the language of sheaves.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: No new theoretical results are included. The examples rely on other well-known results in the studies of dynamical systems and algebraic topology. More complete background is given in the appendices.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The model is described in detail, and the parameters used for figures is included in the captions. While the initial conditions and velocites are not included, the method for choosing the random conditions is mentioned.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code used to create the simulations depicted in the figures is included as supplemental information.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: The parameters used for the simulations are included in the figure captions.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [NA]

   Justification: The paper does not include any experiments for which error bars are relevant.

   Guidelines:

   - The answer NA means that the paper does not include experiments.

- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The simulations preformed, with the small number of agents, did not require significant computational resources. They can be reproduced on any standard modern computer.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The paper conforms to the Code of Ethics by addressing potential harms, of which there are very few.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

Justification: There are no direct pathes to any negative societal impacts that the authors are aware of.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No high risk models were released or proposed.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: To the best of the author's knowledge, the models' original sources are cited. The code included is original to the author. The code is written using Mathematica in the Wolfram Language, a proprietary but common language.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The included code is reasonably well commented and documented.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: No human subjects or crowdsourcing experiments were used.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: No human subjects or crowdsourcing experiments were used.

    Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification: LLMs were not used or relevant for this research.

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.