Detection of functional communities in networks of randomly coupled oscillators using the dynamic-mode decomposition

Christopher W. Curtis*
Department of Mathematics and Statistics, San Diego State University, San Diego, California 92182, USA

Mason A. Porter
Department of Mathematics, University of California, Los Angeles, Los Angeles, California 90095, USA and Santa Fe Institute, Santa Fe, New Mexico 87501, USA

(Received 30 March 2021; accepted 10 August 2021; published 8 October 2021)

Dynamic-mode decomposition (DMD) is a versatile framework for model-free analysis of time series that are generated by dynamical systems. We develop a DMD-based algorithm to investigate the formation of functional communities in networks of coupled, heterogeneous Kuramoto oscillators. In these functional communities, the oscillators in a network have similar dynamics. We consider two common random-graph models (Watts–Strogatz networks and Barabási–Albert networks) with different amounts of heterogeneities among the oscillators. In our computations, we find that membership in a functional community reflects the extent to which there is establishment and sustainment of locking between oscillators. We construct forest graphs that illustrate the complex ways in which the heterogeneous oscillators associate and disassociate with each other.

DOI: 10.1103/PhysRevE.104.044305

I. INTRODUCTION

Researchers in myriad disciplines employ networks to represent entities (i.e., nodes) that interact with each other through their connections (which are encoded by edges) [1]. Network architecture, in turn, affects the dynamics of systems that evolve on networks [2]. Network structure can profoundly impact the spread of diseases [3] and information [4], the collective behavior of coupled oscillators [5,6], and more.

The analysis of networks plays an increasingly important role throughout the sciences and engineering, and this is very prominent in the study of dynamical systems [1,7]. A key question is how interacting entities in a network form collective structures like communities, which can take the form of either structural or functional communities. A structural community is a dense set of nodes that is connected sparsely to other dense sets of nodes [8,9]. A very large number of approaches have been developed to algorithmically detect structural communities in networks. By contrast, considerably less effort has been devoted to the detection of functional communities, which are based on the behaviors or dynamics of the nodes. One way to detect functional communities is by running a dynamical process on a network, constructing a new network (a so-called "functional network") based on the time-series similarity of the outputs of the network's nodes, and then detecting structural communities in the functional network [10,11]. One can also detect functional communities using time-series output of experiments. Whether the time-series output comes from a model or experimental measurements, the focus in functional-community detection is the behavioral similarity of a network's entities over time. These functional communities, whose name is inspired by studies of functional brain networks in neuroscience [11], arise from communities in a network in which the edges encode some type of time-series similarity between the nodes of the network. Another name for such communities is "behavioral communities." Functional communities were explored briefly in [12] using methods from information theory and structural community detection.

In the present study of functional communities, we consider the setting of coupled oscillators on networks. A network itself may have community structure (which is based on its structural communities), which one can examine through one or more of the myriad available methods for community detection [9]. Additionally, by examining the dynamics of coupled oscillators on a network and tracking phenomena such as synchronization, one can also study the same system's functional community structure. This topic was explored by Arenas *et al.* [10] in an investigation of coupled Kuramoto oscillators on networks. The fact that coupled Kuramoto models have been studied so extensively [6] makes them an ideal test case for studies of functional-community detection in coupled oscillators, with a view towards extending such analysis to other systems (including experimental ones, such as in the analysis of neuroimaging data [11]). Although we study coupled oscillators in our work, other types of dynamical systems on networks can also have functional communities [13].

Arenas *et al.* [10] studied Kuramoto oscillators that are coupled to each other on a network with a hierarchical community structure (with smaller, denser communities nested inside larger, sparser ones) and they examined how the architecture of structural communities affects the formation of functional communities, as quantified by how long it takes the oscillators to synchronize. They demonstrated that oscillators in denser communities synchronize faster than oscillators in

*ccurtis@sdsu.edu

sparser ones. Variations of coupled Kuramoto oscillators on networks have also been used to study the coalescence of functional communities [14,15]. In these papers, the network architecture was fixed and the researchers sought to partition the networks based on the dynamics of the oscillators.

In our investigation of functional communities, we take a different perspective from those in the aforementioned papers. We seek to detect communities from output dynamics and we consider both the formation and the disappearance in time of functional communities. To do this, we use the modal-decomposition technique that is known as the dynamic-mode decomposition (DMD). For more information about DMD, see [16–18] for comprehensive reviews and [19–21] for surveys and discussions of recent extensions. The primary benefit of the DMD is that it is a model-free data-processing tool that allows one to generate modal decompositions from arbitrarily complicated data sets over reasonably chosen time intervals. Moreover, in contrast to other modal-decomposition methods (such as principle-component analysis), DMD also gives a convenient mechanism to generate models from measured data alone, and it thereby gives a way to sidestep model development in situations in which it is difficult or even impossible.

In the present paper, instead of examining synchronization timescales as a way to partition a network of oscillators, we generate output data from coupled Kuramoto oscillators on networks (which we construct from random-graph models) and identify functional communities using modes that are generated by the DMD. The flexibility of the DMD allows us to do this in a time-dependent way, and it thereby makes it possible to track the formation and disappearance of communities.

Our DMD-based approach provides a straightforward and flexible way to generate functional communities from a time series of the nodes of a network. In our study of coupled oscillators, the formation of such communities still relies on the synchronization of subsets of the oscillators, but our approach does not require the observation of any global attractors (which do not manifest on the timescales that we examine). This gives an effective method for identifying functional communities in time series of coupled nonlinear oscillators. We anticipate that our approach will be applicable to many classes of time series in which community formation is of interest. Our work complements the recent results of Kunert-Graf *et al.* [22], who used unsupervised learning techniques to cluster DMD modes to capture spatially and temporarily coherent patterns in complex signals. Our code is available at [23].

Our paper proceeds as follows. In Sec. II, we give the necessary background and definitions to understand both DMD and how we generate our data. In Sec. III, we explain our algorithm for the detection of functional communities. In Sec. IV, we present the results of our numerical experiments. In Sec. V, we summarize our results and discuss future work.

II. BACKGROUND AND DEFINITIONS

A. Dynamic-mode decomposition

Consider a nonlinear dynamical system of the form

$$\frac{d\mathbf{y}}{dt} = f(\mathbf{y}, t), \quad \mathbf{y}(0) = \mathbf{x} \in \mathbb{R}^{N_s}. \tag{1}$$

We define the associated flow of (1) by $\mathbf{y}(t) = \varphi(t; \mathbf{x})$. The associated Hilbert space of observables is $L_2(\mathbb{R}^{N_s}, \mathbb{R}, \mu)$, for which we use the shorthand notation $L_2(\mathcal{O})$, where a function $g : \mathbb{R}^{N_s} \longrightarrow \mathbb{R}$ satisfies $g \in L_2(\mathcal{O})$ if

$$\int_{\mathbb{R}^{N_s}} |g(\mathbf{x})|^2 d\mu(\mathbf{x}) < \infty$$

for some appropriate measure μ.

One can gain considerable insight by examining the associated linear representation of the problem. This representation is given by the infinite-dimensional Koopman operator

$$\mathcal{K}^t : L_2(\mathcal{O}) \to L_2(\mathcal{O}).$$

For $g \in L_2(\mathcal{O})$, we have that

$$\mathcal{K}^t g(\mathbf{x}) = g(\varphi(t; \mathbf{x})).$$

The power of moving to a linear-operator framework is that we capture the dynamics of the nonlinear system (1), as measured via observables, using the eigenvalues of $\mathcal{K}^t$. We assume that $\mathcal{K}^t$ has a discrete spectrum. If we can find a basis of $g \in L_2(\mathcal{O})$ via the Koopman eigenfunctions h_j, which satisfy

$$\mathcal{K}^t h_j = e^{t\lambda_j} h_j,$$

then it follows for any other observable g that

$$g = \sum_{j=1}^{\infty} c_j h_j, \quad \mathcal{K}^t g = \sum_{j=1}^{\infty} e^{t\lambda_j} c_j h_j. \tag{2}$$

It is typically impossible to determine the modes of the Koopman operator $\mathcal{K}^t$ in closed form. Therefore, scholars have developed the DMD [16–18] and extended DMD [24,25] for the practical numerical computation of a finite number of Koopman modes.

To do DMD, we start by sampling the flow $\varphi(t; \mathbf{x})$ at discrete times $t_n = t_i + (n-1)\delta t$ (with $n = 1, \ldots, N_T + 1$), where t_i is the initial time and δt is a time step, to generate a data set $\mathbf{y}_n = \varphi(t_n, \mathbf{x})$. If we select the set $\mathbf{g} = \{g_l\}_{l=1}^{M}$ of observables such that $g_l(\mathbf{x}_l) = x_l$ and $M = N_s$, then the DMD approximates $\mathcal{K}^{\delta t}$ by computing the spectra of the finite-dimensional operator $\tilde{K}_a$ that we obtain by solving the optimization problem

$$\mathbf{K}_a = \arg\min_{\mathbf{K}} \|\mathbf{Y}_+ - \mathbf{K}\mathbf{Y}_-\|_F^2,$$

where $\mathbf{K}$ is an $N_s \times N_s$ matrix and

$$\mathbf{Y}_- = [\mathbf{y}_1 \mathbf{y}_2 \cdots \mathbf{y}_{N_T}], \quad \mathbf{Y}_+ = [\mathbf{y}_2 \mathbf{y}_3 \cdots \mathbf{y}_{N_T+1}].$$

After obtaining the $N_s \times N_s$ matrix $\mathbf{K}_a$, we compute its singular-value decomposition $\mathbf{K}_a = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger$ [26], where each factor is an $N_s \times N_s$ matrix, $\mathbf{U}$ and $\mathbf{V}$ are unitary, $\mathbf{\Sigma}$ is diagonal, and the dagger denotes the Hermitian conjugate. Following the standard convention, we order the diagonal entries $\sigma_1, \ldots, \sigma_{N_s}$ (i.e., the singular values) of $\mathbf{\Sigma}$ so that

$$\sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant \sigma_{N_s} \geqslant 0.$$

We then define the diagonal matrix $\tilde{\mathbf{\Sigma}}$ with diagonal entries

$$\tilde{\sigma}_j = \begin{cases} \sigma_j, & \log_{10}(\frac{\sigma_j}{\sigma_1}) \geqslant t_{\mathrm{DMD}} \\ 0, & \log_{10}(\frac{\sigma_j}{\sigma_1}) < t_{\mathrm{DMD}}, \end{cases} \tag{3}$$

where we set the DMD threshold $t_{\rm DMD}$ to remove singular values that are more reflective of ill-conditioning in $\mathbf{K}_a$ than of meaningful information. In practice, we take $t_{\rm DMD} < 0$. We then work with the matrix $\tilde{\mathbf{K}}_a = \mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V}^\dagger$. See [27] for a discussion of this DMD method and related approaches.

We let $\tilde{\mathbf{K}}_a = \mathbf{\Xi} e^{\delta t \mathbf{\Lambda}} \mathbf{\Xi}^{-1}$ and write

$$\mathbf{y}_n = \sum_{j=1}^{N_s} \boldsymbol{\xi}_j e^{n\,\delta t\,\lambda_j} h_j(\mathbf{x}), \quad h_j(\mathbf{x}) = (\mathbf{\Xi}^{-1}\mathbf{x})_j .$$

The real part of the eigenvalue λ_j gives the amplitude of the jth mode, and the imaginary part of λ_j gives its oscillation frequency. We determine the evolution of the associated eigenfunctions of the Koopman operator using the formula

$$h_j(\mathbf{y}_n) = (\mathbf{\Xi}^{-1}\mathbf{y}_n)_j .$$

One measure of error is the extent to which the computed modes $h_j(\mathbf{x})$ behave as Koopman eigenfunctions [21]. To quantify this, we follow [28] and calculate

$$\mathcal{E}_j = \frac{\sum_{n=1}^{N_T} |h_j(\mathbf{y}_{n+1}) - e^{\delta t \lambda_j} h_j(\mathbf{y}_n)|}{\sum_{n=1}^{N_T} |h_j(\mathbf{y}_n)|}$$

for each $j \in \{1, \ldots, N_s\}$. The quantity $\mathcal{E}_j$ gives a normalized measure of how well the computed approximations to the Koopman eigenfunctions and eigenvalues are able to linearize the dynamics as in (2). We choose a tolerance ϵ_m and keep only the modes that satisfy $\mathcal{E}_j < \epsilon_m$. We denote the number of the modes that we keep by N_r. We enforce how well these N_r modes reconstruct the time series by choosing a reconstruction-error parameter $\epsilon_{\rm rc}$ so that

$$\frac{\|\mathbf{Y}_+ - \mathbf{H}\|_F}{\|\mathbf{Y}_+\|_F} < \epsilon_{\rm rc} ,$$

where $\mathbf{H}$ is the matrix with columns

$$\mathbf{h}_n = \sum_{l=1}^{N_r} \boldsymbol{\xi}_{j_l} e^{n\,\delta t\,\lambda_{j_l}} h_{j_l}(\mathbf{x})$$

and j_l indexes the subset of modes that satisfy the criterion $\mathcal{E}_j < \epsilon_m$.

There is an interplay between the choices of $t_{\rm DMD}$, ϵ_m, and $\epsilon_{\rm rc}$. It takes effort (e.g., through trial and error) to balance these parameters to produce meaningful results. For example, if $t_{\rm DMD}$ is too small (e.g., $t_{\rm DMD} = -16$, which corresponds to machine precision on most desktop computers), then one typically corrupts a DMD computation to the point that there is no practical way to find reasonable choices of ϵ_m or $\epsilon_{\rm rc}$. However, if $t_{\rm DMD}$ is too large, only a very small number of modes satisfy the $\mathcal{E}_j < \epsilon_m$ criterion, rendering it difficult to make reasonable choices of $\epsilon_{\rm rc}$. Likewise, setting ϵ_m to a value that is too small can produce excellent approximations to Koopman modes, but it is hard for the modal reconstruction to allow reasonable choices of $\epsilon_{\rm rc}$. In Sec. IV, we describe parameter choices that reflect the necessary balancing.

B. Coupled oscillators on random graphs

We describe a graph (i.e., network) using an associated adjacency matrix $\mathbf{A}$ with elements A_{jk}. We assume that each of our graphs, which we assemble using random-graph models, is undirected. Therefore, their associated adjacency matrices are symmetric. We also assume that each graph has no self-edges (so $A_{jj} = 0$ for all j) and no multiedges. Finally, we assume that our graphs are unweighted, so all entries of each matrix $\mathbf{A}$ are either 1 or 0.

To study dynamics on a graph, suppose that each node $j \in \{1, \ldots, N_s\}$ is associated with a Kuramoto oscillator [6]. This yields the dynamical system

$$\dot{\theta}_j = \omega_j + \frac{K}{N_s} \sum_{k=1}^{N_s} A_{jk} \sin(\theta_k - \theta_j), \quad \omega_j \sim \frac{1}{\gamma} \tilde{g}\left(\frac{x}{\gamma}\right), \quad (4)$$

where $\theta_j \in [0, 2\pi)$ is the phase of the jth oscillator, ω_j is the natural frequency of the jth oscillator, $K \geqslant 0$ controls the coupling strength between oscillators, and $\tilde{g}(y)$ is a probability distribution with mean 0 and width 1 (which constitutes a variance of 1, provided the variance of the distribution $\tilde{g}$ is well-defined). Because of our rescaling, the parameter γ is the variance of the distribution; we obtain identical oscillators in the limit $\gamma \to 0^+$.

We are interested in the extent to which the oscillators lock. In the strongest sense, locking means phase locking, which is defined as

$$\lim_{t\to\infty} \theta_j(t) = \theta_p, \quad j \in \{1, \ldots, N_s\}. \quad (5)$$

We also consider frequency locking, which is defined as

$$\lim_{t\to\infty} \dot{\theta}_j(t) = \omega_f, \quad j \in \{1, \ldots, N_s\}.$$

In typical scenarios, phase locking implies frequency locking (although this need not be true in the presence of noise [29]), but the converse is not true in general.

To measure the extent that the oscillators lock, we calculate the order parameter

$$r_p(t) e^{i\psi_p(t)} = \frac{1}{N_s} \sum_{j=1}^{N_s} e^{i\theta_j(t)},$$

and we note that $0 \leqslant r_p(t) \leqslant 1$. If the oscillators are equally spaced at time t, such that $\theta_j = \frac{2\pi(j-1)}{N_s}$, then

$$r_p(t) e^{i\psi_p(t)} = \frac{1}{N_s} \sum_{j=1}^{N_s} (e^{2\pi i/N_s})^{(j-1)} = 0,$$

so a value of $r_p(t)$ that is sufficiently close to 0 for finite N_s indicates an absence of phase locking between the oscillators.

If we satisfy the phase-locking criterion (5), then $r_p(t) \to 1$ and

$$r_p(t) e^{i\psi_p(t)} \to e^{i\theta_p} \quad \text{as } t \to \infty .$$

For frequency locking, it can be true that

$$\theta_j(t) \to \omega_f t + \theta_{j,s} \quad \text{as } t \to \infty ,$$

where $\theta_{j,s}$ is a phase shift. In this case,

$$r_p(t) e^{i\psi_p(t)} \to e^{i\omega_f t} \frac{1}{N_s} \sum_{j=1}^{N_s} e^{i\theta_{j,s}} \quad \text{as } t \to \infty . \quad (6)$$

Depending on the particular locations of the phase shifts $\theta_{j,s}$, it is possible that $r_p(t)$ is small in magnitude. However, as one can see from Eq. (6), we also expect $r_p(t)$ to be almost constant in time. We use practical criteria to determine when oscillators are phase locked and/or frequency locked. We understand a collection of oscillators as being close to a phase-locked state when $r_p(t)$ is close to 1 in magnitude and does not vary much in time. We understand a collection of oscillators as being close to a frequency-locked state if $r_p(t)$ is not too close to 0 and does not vary much in time.

Because we are characterizing the closeness to locking in an imprecise way, we do perturbation theory to facilitate the interpretation of our later numerical results. We perturb with respect to the mean $\langle r_p \rangle$ of $r_p(t)$ over the time interval $[t_i, t_f]$. To compute this mean, we calculate

$$\langle r_p \rangle = \frac{1}{t_f - t_i} \int_{t_i}^{t_f} r_p(t) dt \, .$$

For frequency locking, a natural condition for our perturbations is to suppose as t becomes large that we can write each phase $\theta_j(t)$ in the form

$$\theta_j(t) = \omega_f t + \tilde{\theta}_j(\epsilon t) \, , \quad 0 < \epsilon \ll 1 \, . \tag{7}$$

The ansatz in Eq. (7) allows slow modulations around the locking frequency, where we control the extent of the slowness using the small positive parameter ϵ. We see that

$$r_p(t) = \left| \frac{1}{N_s} \sum_{j=1}^{N_s} e^{i\theta_j(t)} \right| = \left| \frac{1}{N_s} \sum_{j=1}^{N_s} e^{i\tilde{\theta}_j(\epsilon t)} \right| ,$$

which demonstrates that $r_p(t)$ varies slowly with respect to time. We then write the order parameter as $r_p(t) = \tilde{r}_p(\epsilon t)$, where

$$\tilde{r}_p(\epsilon t) = \left| \frac{1}{N_s} \sum_{j=1}^{N_s} e^{i\tilde{\theta}_j(\epsilon t)} \right| .$$

Making the reasonable assumption that $\tilde{r}_p(\epsilon t)$ has a well-defined derivative $\dot{\tilde{r}}_p(\epsilon t)$, we then compare $r_p(t)$ to its mean value for $t \in [t_i, t_f]$ and write

$$\begin{aligned} r_p(t) - \langle r_p \rangle &= \frac{1}{t_f - t_i} \int_{t_i}^{t_f} [\tilde{r}_p(\epsilon t) - \tilde{r}_p(\epsilon s)] ds \\ &= \frac{1}{t_f - t_i} \int_{t_i}^{t_f} [\tilde{r}_p(\epsilon s + \epsilon(t - s)) - \tilde{r}_p(\epsilon s)] ds \\ &= \frac{\epsilon \dot{\tilde{r}}_p(\epsilon \chi)}{t_f - t_i} \int_{t_i}^{t_f} (t - s) ds \quad [\text{with } \chi \in (t_i, t_f)] \\ &= \dot{\tilde{r}}_p(\epsilon \chi) \left(\epsilon t - \epsilon \frac{t_i + t_f}{2} \right) , \end{aligned} \tag{8}$$

where we have used the mean-value theorem for the penultimate equality. This calculation shows that if oscillators vary slowly around a frequency-locked state, then the variation of the magnitude of the order parameter around the mean should (i) also be slow and (ii) be bounded by the magnitude of the derivative multiplied by the scaled difference between the time and the midpoint of the time interval over which we average. As our numerical results demonstrate (see Sec. IV), frequency locking is more typical than phase locking in our networks of heterogeneous oscillators. Therefore, throughout the remainder of our paper, we interpret our results in the context of Eq. (8).

Similarly, we claim that a reasonable perturbation around a phase-locked state is

$$\theta_j(t) = \theta_p + \epsilon \tilde{\theta}_j(t) \, , \quad 0 < \epsilon \ll 1 \, .$$

Using the Taylor-series expansion

$$e^{i\epsilon\tilde{\theta}_j} \approx 1 + i\epsilon\tilde{\theta}_j - \frac{\epsilon^2}{2}\tilde{\theta}_j^2$$

yields

$$\begin{aligned} r_p(t) &\approx \left[1 + \epsilon^2 \left(\tilde{m}^2(t) - \frac{1}{N_s} \sum_{j=1}^{N_s} \tilde{\theta}_j^2(t) \right) \right]^{1/2} \\ &\approx 1 + \frac{\epsilon^2}{2} \left(\tilde{m}^2(t) - \frac{1}{N_s} \sum_{j=1}^{N_s} \tilde{\theta}_j^2(t) \right) \\ &\approx 1 - \frac{\epsilon^2}{2N_s} \sum_{j=1}^{N_s} [\tilde{\theta}_j(t) - \tilde{m}(t)]^2 \, , \end{aligned} \tag{9}$$

where

$$\tilde{m}(t) = \frac{1}{N_s} \sum_{j=1}^{N_s} \tilde{\theta}_j(t) \, .$$

In the second line of (9), we used the Taylor-series approximation $\sqrt{1 + x^2} \approx 1 + x^2/2$. From Eq. (9), we can now see clearly how variances around the mean of the perturbations of the locked phase θ_p reduce the magnitude of the order parameter $r_p(t)$ below 1 and introduce temporal variations. We also anticipate that if the perturbations $\tilde{\theta}_j$ vary slowly, then $r_p(t)$ should also vary slowly.

The issues that arise from the above choice of order parameter have inspired efforts, such as the use of topological data analysis [30], to apply data-analysis techniques to study order parameters in complex systems. Moreover, in addition to the global order parameter that we employ, one can calculate other order parameters to examine localized locking in subsets of oscillators [31–33]. Applying a DMD approach with these alternative order parameters is worthwhile to explore in future research efforts.

Additionally, as was shown in [34], one can derive analytical criteria that guarantee phase locking of some subset of Kuramoto oscillators on a network. Specifically, in the asymptotic limit of infinitely many oscillators and arbitrarily long times, there is necessarily some locking of some subset of the oscillators if

$$\gamma < \gamma_c \, , \quad \gamma_c = \frac{\pi N_s K \tilde{g}(0) \sigma_{\max}^{\mathrm{A}}}{2} \, ,$$

where $\sigma_{\max}^{\mathrm{A}}$ is the largest eigenvalue of the adjacency matrix $\mathbf{A}$. However, the number of oscillators that lock (and the number of communities of oscillators that lock) is not determined by this criterion. Moreover, in the present paper, we consider systems with a relatively small number of oscillators, so we

are far away from the asymptotic regime that was studied in [34].

III. ESTIMATION OF LOCKING AND FUNCTIONAL-COMMUNITY DETECTION USING THE DYNAMIC-MODE DECOMPOSITION

To measure the amount of locking among the oscillators in the time interval $[t_i, t_f]$, we examine the collection $\{\boldsymbol{\xi}_m\}_{m=1}^{N_r}$ of Koopman modes. We define the overlap matrix $\mathbf{C}^{\mathrm{o}}$ with elements

$$C^{\mathrm{o}}_{jl} = \left| \sum_{m=1}^{N_r} \hat{\xi}_{m,j} \hat{\xi}^*_{m,l} \right|, \quad \hat{\boldsymbol{\xi}}_m = \frac{\boldsymbol{\xi}_m}{\|\boldsymbol{\xi}_m\|_2}.$$

Using the Cauchy–Schwarz inequality, we know that

$$0 \leqslant C^{\mathrm{o}}_{jl} \leqslant 1.$$

We also see that $\mathbf{C}^{\mathrm{o}}$ is symmetric. Using a threshold $C_{\mathrm{cr}} \in (0, 1)$, we construct graphs of the strongest interactions by generating an adjacency matrix $\mathbf{A}^{(\mathrm{md})}(t_i, t_f)$ with elements

$$A^{(\mathrm{md})}_{jl}(t_i, t_f) = \begin{cases} 1, & C^{\mathrm{o}}_{jl} \geqslant C_{\mathrm{cr}} \\ 0, & C^{\mathrm{o}}_{jl} < C_{\mathrm{cr}}. \end{cases}$$

Because $\mathbf{C}^{\mathrm{o}}$ is symmetric, it is also true that $\mathbf{A}^{(\mathrm{md})}(t_i, t_f)$ is symmetric and thus that its associated graph is undirected. In a given time interval, we say that subsets of the oscillators belong to a community if they are a part of the same connected component of this graph.

We expand on this notion of community to track the merging, splitting, formation, and dissolution of communities over time. This adds to a growing body of literature on detecting communities in time-varying networks; see, for example, [11,35–37]. We separate a fixed time interval $[t_i, t_f]$ into n_w subintervals $\{I_j\}_{j=1}^{n_w}$ of equal length. We also generate $n_w - 1$ equally spaced intervals $\{\tilde{I}_j\}_{j=1}^{n_w-1}$ that connect the midpoints of the I_j intervals. We order these intervals to give a total of $2n_w - 1$ intervals $\bar{I}_j$ such that

$$\bar{I}_j = \begin{cases} I_{(j+1)/2}, & 1 \equiv j \bmod 2 \\ \tilde{I}_{j/2}, & 0 \equiv j \bmod 2. \end{cases}$$

With this construction, our collection of intervals overlaps in time. We generate a sequence of communities using the sequence

$$\{\mathbf{A}^{(\mathrm{md})}(\bar{I}_j)\}_{j=1}^{2n_w-1}$$

of adjacency matrices. We use overlapping time intervals to help ensure some continuity of the communities from one interval to the next. See [38] for further exploration of this issue.

To study how communities evolve over time, we define a time-dependent graph $\mathcal{G}_{t_i,t_f}$ that tracks community relationships across time. To generate this graph, for the first time interval $\bar{I}_1$, we use $\mathbf{A}^{(\mathrm{md})}(\bar{I}_1)$ to obtain communities $\{C_k(\bar{I}_1)\}_{k=1}^{N_{c,1}}$ of oscillators in which each community corresponds to one connected component of the graph that is associated with $\mathbf{A}^{(\mathrm{md})}(\bar{I}_1)$. Each of these communities corresponds to an individual node in $\mathcal{G}_{t_i,t_f}$. After initializing $\mathcal{G}_{t_i,t_f}$, we generate successive graphs in a sequence using the following recursive process. Assuming that we have constructed $\mathcal{G}_{t_i,t_f}$ up to the jth interval, we place an edge between the kth community $C_k(\bar{I}_j)$ from $\mathbf{A}^{(\mathrm{md})}(\bar{I}_j)$ and the lth community $C_l(\bar{I}_{j+1})$ from $\mathbf{A}^{(\mathrm{md})}(\bar{I}_{j+1})$ if the nodes in $C_k(\bar{I}_j)$ are identical to or are a subset of those in $C_l(\bar{I}_{j+1})$. Otherwise, if the nodes in $C_k(\bar{I}_j)$ are not contained within a community or have fractured into smaller communities at the $(j+1)$th interval, then $C_k(\bar{I}_j)$ terminates at the jth interval. This process ensures that $\mathcal{G}_{t_i,t_f}$ is a forest that adaptively connects communities to each other across time based on the relative overlaps of Koopman modes. In our numerical computations (see Sec. IV), we observe coalescence of communities when there is phase locking or frequency locking. We summarize our community-detection algorithm to generate community-relationship graphs $\mathcal{G}_{t_i,t_f}$ in Algorithm 1.

Algorithm 1. Thresholding algorithm for functional-community detection.

Given a time series $\{\mathbf{y}_n\}_{n=1}^{N_s}$ in the time interval $[t_i, t_f]$
Require: Choose C_{cr} and n_w. Separate the interval $[t_i, t_f]$ into overlapping subintervals $\{\bar{I}_j\}_{j=1}^{2n_w-1}$.
Output: Community-relationship graph $\mathcal{G}_{t_i,t_f}$
for $j \in \{1, \ldots, 2n_w - 1\}$ **do**
1. Compute the DMD of the portion of the time series that is in the interval $\bar{I}_j$.
2. Compute the overlap matrix $\mathbf{C}^{\mathrm{o}}$ and its associated adjacency matrix $A^{(\mathrm{md})}(\bar{I}_j)$ using the threshold C_{cr}.
3. Compute the associated communities $\{C_k(\bar{I}_j)\}_{k=1}^{N_{c,j}}$. For each of the $N_{c,j}$ communities, add a node to $\mathcal{G}_{t_i,t_f}$.

if $j > 1$ **then**
for $k \in \{1, \ldots, N_{c,j}\}$ **do**
for $l \in \{1, \ldots, N_{c,j-1}\}$ **do**
if $C_k(\bar{I}_j) \subseteq C_l(\bar{I}_{j-1})$ **then**
Form an edge between the corresponding nodes in $\mathcal{G}_{t_i,t_f}$.
Break
end if
end for
end for
end if
end for

IV. NUMERICAL EXPERIMENTS

We now conduct a variety of numerical experiments to examine how successfully DMD can detect functional communities in networks of Kuramoto oscillators. Our networks have $N_s = 800$ oscillators, and the coupling parameter is $K = 10$. We examine networks with as many as $N_s = 3200$ oscillators and we obtain results (not shown) that are similar to the ones that we discuss below. We draw the natural frequencies of the oscillators from a Gaussian distribution, so $\gamma_c = 4000\sqrt{\pi}\sigma^{\mathrm{A}}_{\mathrm{max}}$. Given this choice of distribution, we expect that at least a subset of the oscillators will become at least nearly locked because $\sigma^{\mathrm{A}}_{\mathrm{max}} > 1$ in all of our experiments. We use a second-order Runge–Kutta solver with a step size of $\delta t = 0.05$ to integrate the Kuramoto model (4). We run

all simulations up to $t_f = 800$. From $t = 760$ to $t_f = 800$, we perform DMD on the circle around which the oscillators move, so we use the coordinates $(\cos[\theta_j(t)], \sin[\theta_j(t)])$.

We set the DMD threshold in Eq. (3) to $t_{\text{DMD}} = -4$, the error threshold of the DMD modes to $\epsilon_m = 10^{-1}$, and the reconstruction error to $\epsilon_{\text{rc}} = 10^{-1}$. As we described at the end of Sec. II A, these parameter choices reflect a balance that we obtain through trial and error. In our numerical experiments, we report results for heterogeneity parameters $\gamma = 0.1$, $\gamma = 1$, and $\gamma = 10$. Although smaller values of γ (indicating more homogeneous oscillators) allow wider viable ranges for the three DMD parameters, our choices work well across the two-decade range of the values of γ.

In our simulations, we use two classical random-graph models (RGMs): the Barabási–Albert (BA) preferential-attachment model [39] and the Watts–Strogatz (WS) small-world model [40]. After specifying the parameter values of an RGM, we generate a single network from the model. When we simulate the Kuramoto oscillators on either network, we initialize our dynamical system with N_s oscillators that we space uniformly around a circle. In the following paragraphs, we indicate the specific variants and parameter values that we discuss for these models. We also examined other parameter choices for these RGMS, and we obtained similar results.

For our BA network, we start with a seed network of $n = 10$ nodes with edges that we choose according to a $G(n, p)$ Erdős–Rényi model with an independent, uniform probability of $p = 0.75$ for placing an edge between nodes. At each time step, we add a new node to the network; this node has four edges that we connect uniformly at random (without replacement) to existing nodes. We grow our BA network until it consists of $N_s = 800$ nodes.

Our WS network also has $N_s = 800$ nodes. We start with a ring lattice in which each node is adjacent to $K_{\text{nbh}} = 20$ nearest neighbors. We then rewire the network as follows. Starting at the jth node n_j, we consider half of its nearest-neighbor nodes, where the specific neighbors n_l are those with the indices

$$l = \tilde{l} \bmod N_s\,, \quad j < \tilde{l} \leqslant j + \frac{K_{\text{nbh}}}{2}\,.$$

For each n_l, with probability β, we do the following:

(i) We choose a node n_k uniformly at random from the nodes that are not adjacent to n_j.

(ii) We remove the edge e_{jl} that connects n_j to n_l.

(iii) We add an undirected edge e_{jk} to connect n_j to n_k.

As we iterate through the nearest neighbors of n_j, we do not allow rewiring to any nodes that currently or formerly are nearest neighbors. However, such a nearest-neighbor connection can arise again when we rewire connections from other nodes after we have moved on from n_j. Therefore, it is possible for nearest-neighbor edges to be removed and added back again, although this is not common. In our simulations, we set the rewiring probability to $\beta = 0.6$. As discussed in [1], the rewiring process drastically reduces the mean geodesic distance between nodes from that of the initial ring lattice.

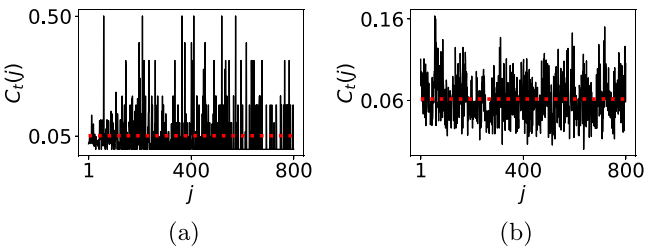

FIG. 1. Local clustering coefficient $C_t(j)$ of each node j in (a) a Barabási–Albert network and (b) a Watts–Strogatz network. The horizontal line indicates the mean value of $C_t(j)$. Neither of these networks tends to have particularly large clustering coefficients, although the mean local clustering coefficient is larger in the WS network than in the BA network because of the nature of the associated RGMs.

In each of our examples, we study the structure of the largest connected component (LCC) of the graph[1] that is associated with the adjacency matrix $\mathbf{A}^{(\text{md})}([760, 800])$. Our method of detecting functional communities always generates some output (regardless of whether it is interpretable). It is instructive to examine the details of the predominant community that we obtain in the final examined time interval. We thereby improve our understanding of the meaning of the communities that we obtain using our approach. To help illustrate the structure of the LCC, we compare the local clustering coefficients of the LCC to those in the original RGM. Note that all of our subsequent calculations with the oscillators in the LCC use the time-series output of our original simulation. This yields our order parameters.[2]

For the jth node of a network, we calculate the local clustering coefficient [1]

$$C_t(j) = \frac{\text{number of triangles through node } j}{\binom{d_j}{2}}\,,$$

where $\binom{d_j}{2}$ denotes the number of ways to choose 2 of d_j items and d_j denotes the degree of node j. In Fig. 1, we plot the local clustering coefficients of each node in networks that we construct using the BA and WS models. For most nodes, $C_t(j)$ is not particularly large for either RGM. Therefore, we do not expect to observe particularly modular structural communities in either RGM. In the WS model, the local clustering coefficients tend to be small because the rewiring probability is large.

To generate communities such as the LCC, we set the threshold of the correlation graph to be $C_{\text{cr}} = 0.99$, except for the WS network with heterogeneity parameter $\gamma = 10$, where

[1]For expository convenience, we typically describe these LCCs as properties of the adjacency matrices themselves. We sometimes also refer to these adjacency matrices as networks.

[2]We consider all of the oscillators when computing these order parameters. In particular, we do not compute the order parameters by simulating Kuramoto dynamics using only oscillators that are in the LCC.

we use $C_{\rm cr} = 0.98$. We do this because the threshold $C_{\rm cr} = 0.99$ only yields one or two oscillators in the LCC, which makes it difficult to measure locking. Using a slightly smaller value of $C_{\rm cr}$ in this case ameliorates this issue and yields more intelligible results. Among the values of $C_{\rm cr}$ that we examine, these parameter choices typically produce the best community discrimination across the examined values of γ. Smaller values of $C_{\rm cr}$ generate fewer (and larger) communities. At large values of γ, this can allow too much heterogeneity within communities, and we then observe little to no locking. For example, using $C_{\rm cr} = 0.95$ yields very similar results as $C_{\rm cr} = 0.99$ when $\gamma = 0.1$, but it yields communities with much less locking than for $C_{\rm cr} = 0.99$ when $\gamma = 1$ or $\gamma = 10$ in the BA model. Our observations for the WS model are similar. Arguably, it may be desirable to allow the value of $C_{\rm cr}$ to change to adjust to the known heterogeneity (as quantified by γ) of the oscillators and produce communities with the most locking, but this increases the difficulty of comparing the functional communities that we obtain for different values of γ. Therefore, we use the value $C_{\rm cr} = 0.99$ for most of our case studies to facilitate exposition and consistency in reporting results.

Because we know about a key measurable aspect (which, in our case, is the amount of locking of the oscillators) of the underlying model, we are able to adjust our choices of parameter values to account for it. We expect our approach to be useful for problems in which the dynamics includes a similarly recognizable and quantifiable feature. There are many ways to define and identify such features in what one can call the task of "coherent-structure identification." (See [41–43] for discussions of coherent structures in fluids and other dynamical systems.) Without such features, our DMD approach is likely to produce arbitrary results that are difficult to analyze and verify, and we would not suggest using our method in such situations.

It is also worth considering other RGMs and multiple instances of an RGM with the same parameter values. We performed calculations using stochastic block models [1] and obtained results that are similar to those that we report in the present paper. If one considers multiple instances of an RGM, one can average over the $\mathbf{C}^{\rm o}$ matrices to obtain a consensus matrix and then do community detection using the consensus matrix. This is an interesting question that is worth exploring in future work.

A. Weakly heterogeneous oscillators: $\gamma = 0.1$

We first consider oscillators that are only weakly heterogeneous by examining the case $\gamma = 0.1$. As one can see in Figs. 2(a) and 2(b), in which we compare the distribution of oscillator frequencies $\dot{\theta}_j$ at $t = 0$ and $t = t_f = 800$, the mostly homogeneous natural frequencies of the oscillators tends to yield frequency locking. In these figures, we also observe differences that arise from the different network topologies of WS and BA networks. There is only a small amount of locking in the BA network, whereas the WS network is in an almost fully locked state. The similarity in the dynamics of the oscillators and the relatively long simulation times in both networks results in a relatively modest number of DMD

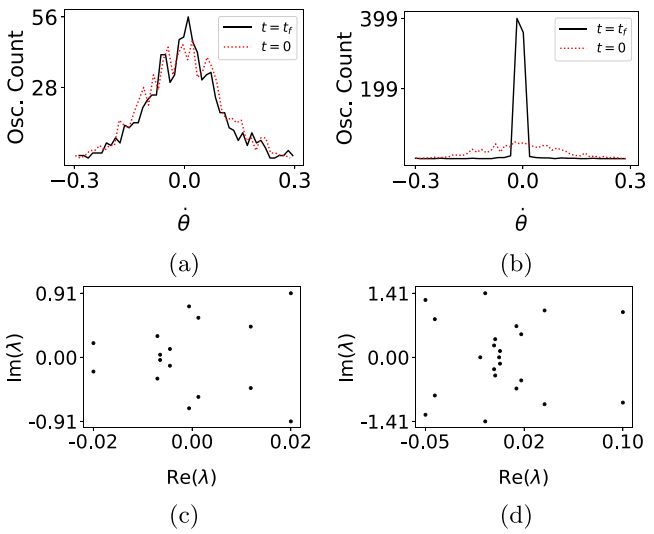

FIG. 2. The (a,b) frequency distributions and (c,d) DMD spectra for weakly heterogeneous Kuramoto oscillators (i.e., when $\gamma = 0.1$). We show our results for (a,c) a BA network and (b,d) a WS network. In (a) and (b), we plot the distributions of $\dot{\theta}_j$ at $t = 0$ and $t = t_f$. Observe that the WS network clearly exhibits locking, whereas the BA network does not have a significant shift in the initial distribution of oscillator frequencies. In (c) and (d), we show the associated DMDs of these networks. The BA network has a smaller frequency range than the WS network, as one can see in the values of Im(λ). Additionally, the larger range of Re(λ) values in the WS network than in the BA network indicates that the former experiences larger changes in DMD-mode amplitudes. In (a) and (b), "Osc." stands for "Oscillator."

modes [see Figs. 2(c) and 2(d)], with only about ten modes in each case.

In Figs. 3(a) and 3(b), we show the sizes of the communities that we generate with a given correlation-graph threshold. In both the BA and WS networks, the LCC of $\mathbf{A}^{(\rm md)}$ is the largest community by far and almost all other oscillators belong to their own single-oscillator communities. As expected from the locking dynamics, the LCC is much smaller in the BA network than in the WS network. In Figs. 3(c) and 3(d), we see that the mean local clustering coefficient in the BA network is smaller than that in the WS network. However, in both networks, the values of the local clustering coefficient are markedly larger in the LCC of $\mathbf{A}^{(\rm md)}$ than in the original adjacency matrices of these networks. Compare Figs. 3(c) and 3(d) to Figs. 1(a) and 1(b).

In Fig. 4, we plot the order parameters both for the original BA and WS networks and for the associated LCCs of $\mathbf{A}^{(\rm md)}$. We see in Figs. 4(e) and 4(f) (which we obtain after applying the thresholding from Algorithm 1) that the order-parameter magnitude $r_p(t)$ oscillates less around the mean than in Figs. 4(a) and 4(b). By comparing the real parts of the phases of the order parameters in Figs. 4(c) and 4(g), we see that the LCC of the BA network has stronger frequency locking than the original network. By contrast, given that the original WS network is almost fully phase locked, its LCC only exhibits a little bit more phase locking [see Figs. 4(d) and 4(h)].

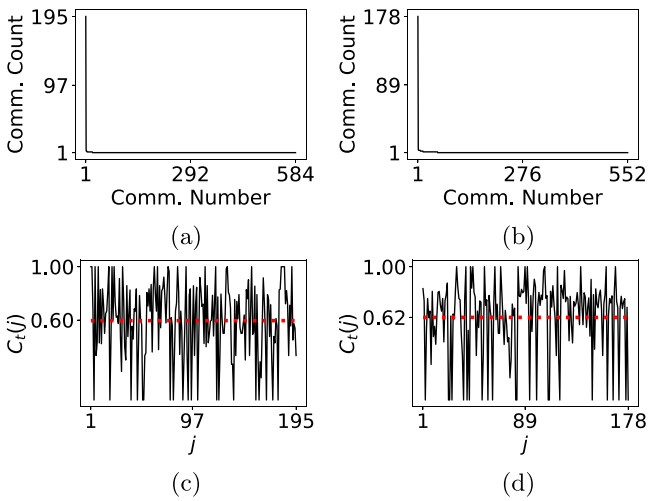

FIG. 3. The (a,b) numbers of nodes in each connected component of $\mathbf{A}^{(\mathrm{md})}$ and the (c,d) local clustering coefficients of the nodes in the LCC of $\mathbf{A}^{(\mathrm{md})}$ of the BA and WS networks for oscillators with $\gamma = 0.1$. The horizontal line in (c) and (d) indicates the mean value of the local clustering coefficient. We show our results for (a,c) a BA network and (b,d) a WS network. In both the BA network and the WS network, the LCC is the largest community by far. Additionally, in both networks, the local clustering coefficients of the nodes in the LCC of $\mathbf{A}^{(\mathrm{md})}$ are markedly larger than the baseline values in the original networks in Fig. 1. In (a) and (b), "Comm." stands for "Community."

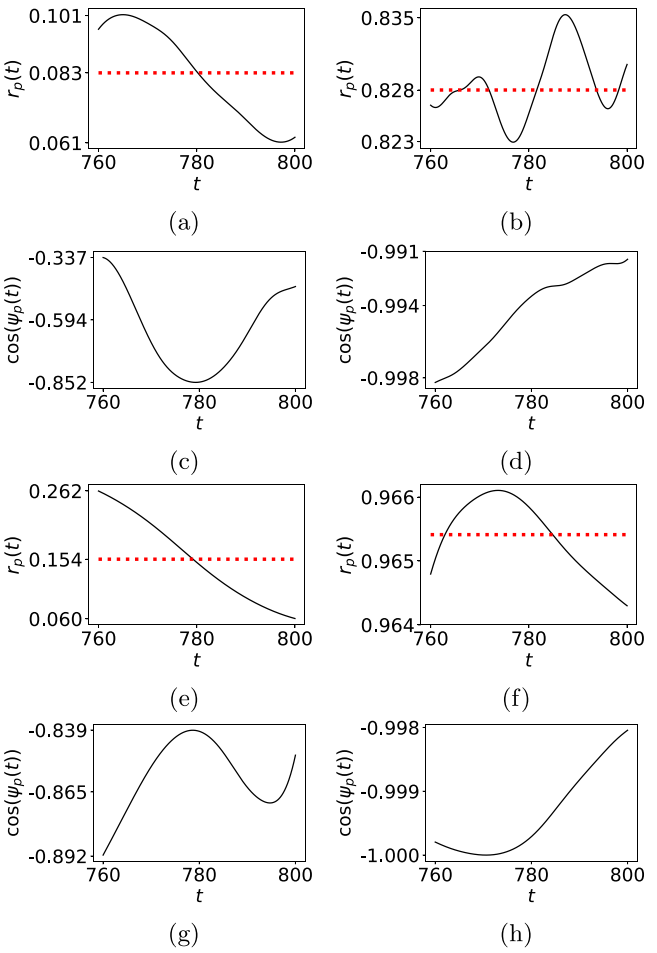

FIG. 4. The (a,b,e,f) order-parameter magnitude $r_p(t)$ and (c,d,g,h) real part $\cos[\psi_p(t)]$ of the order-parameter phase for the BA and WS networks with weakly heterogeneous oscillators (i.e., when $\gamma = 0.1$). We show our results for (a,c,e,g) a BA network and (b,d,f,h) a WS network. Our results for the original two networks are in (a)–(d) and our results for the LCCs of $\mathbf{A}^{(\mathrm{md})}$ of these networks are in (e)–(h). The horizontal line in (a), (b), (e), and (f) indicates the mean $\langle r_p(\cdot)\rangle$ of the order-parameter magnitude. In the BA network, we observe stronger frequency locking in the LCC of $\mathbf{A}^{(\mathrm{md})}$ than in the original network because the former has slower variations in $r_p(t)$ around the mean and a phase that behaves more regularly. In the WS network, we see evidence that the LCC of $\mathbf{A}^{(\mathrm{md})}$ has more phase locking than in the original network.

B. Moderately heterogeneous oscillators: $\gamma = 1$

As we anticipated and confirmed in our numerical computations, the small variance in the distribution of frequencies ω_j when $\gamma = 0.1$ tends to result in locked states. The situation is starkly different when $\gamma = 1$ and the natural frequencies of the oscillators are thus much more heterogeneous. As we can see in Figs. 5(a) and 5(b), the system does not appear to reach a locked state on the same timescales that we observed when $\gamma = 0.1$. Additionally, the DMD spectra in Figs. 5(c) and 5(d) now have about 50 modes, indicating a far greater complexity in the dynamics when $\gamma = 1$ than when $\gamma = 0.1$. This, in turn, results in much smaller functional communities [see Figs. 6(a) and 6(b)]. Nevertheless, for both $\gamma = 0.1$ and $\gamma = 1$, the local clustering coefficients of the LCC of the $\mathbf{A}^{(\mathrm{md})}$ networks are much larger than those in the original networks.

This clear absence of locking is echoed in Figs. 7(a) and 7(b), which reveal that the order-parameter magnitudes $r_p(t)$ are markedly more oscillatory when $\gamma = 1$ than they are when $\gamma = 0.1$. Additionally, we do not observe steady oscillations in $\cos[\psi_p(t)]$; see Figs. 7(c) and 7(d). We see in Figs. 7(e) and 7(f) that in the LCC of $\mathbf{A}^{(\mathrm{md})}$, the oscillations of $r_p(t)$ around the mean are slower than those in the original networks. Using Eq. (8), we attribute this observation to the associated much stronger frequency locking of the oscillators in the LCC. The far steadier and more uniformly oscillating phases in Figs. 7(g) and 7(h) than in Figs. 7(c) and 7(d) corroborate this observation.

C. Strongly heterogeneous oscillators: $\gamma = 10$

Finally, we examine strongly heterogeneous oscillators by setting $\gamma = 10$. Recall that we use a threshold of $C_{\mathrm{cr}} = 0.99$ for the BA model and a threshold of $C_{\mathrm{cr}} = 0.98$ for the WS model. We use a smaller threshold for the WS model to ensure that there are enough oscillators in the LCC to obtain meaningful measurements. Based on our previous results, we anticipate that we are unlikely to observe locking in any significant fraction of the oscillators on the timescales that we examine. Our computations in Fig. 8 confirm this expectation.

In Figs. 8(a) and 8(b), we see that there is almost no change in the frequency distributions during our simulations. This

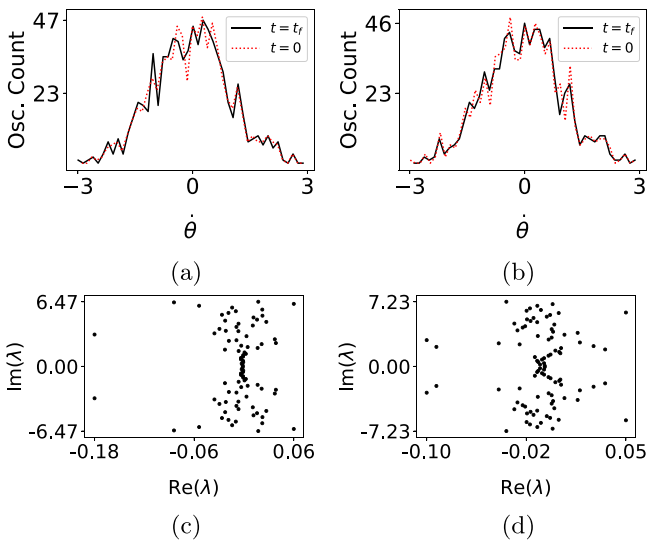

FIG. 5. The (a,b) frequency distributions and (c,d) DMD spectra for moderately heterogeneous Kuramoto oscillators (i.e., when $\gamma = 1$). We show our results for (a,c) a BA network and (b,d) a WS network. We plot the distributions of $\dot{\theta}_j$ at $t = 0$ and $t = t_f$ for (a) a BA network and (b) a WS network. When $\gamma = 1$, the characteristics of the networks from the two RGMs are much less distinguishable than when $\gamma = 0.1$. However, the DMD spectra still have some differences, such as how eigenvalues cluster on the imaginary axis. In (a) and (b), "Osc." stands for "Oscillator."

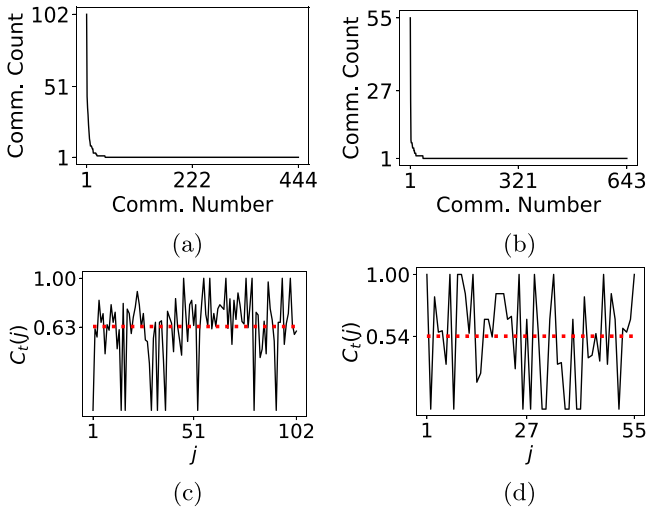

FIG. 6. The (a,b) numbers of nodes in each connected component of $\mathbf{A}^{(\mathrm{md})}$ and the (c,d) local clustering coefficients of the nodes in the LCC of $\mathbf{A}^{(\mathrm{md})}$ of the BA and WS networks for oscillators with $\gamma = 1$. The horizontal line in (c) and (d) indicates the mean value of the local clustering coefficient. We show our results for (a,c) a BA network and (b,d) a WS network. The LCCs are still markedly larger than any other community, but they are much smaller than they are when $\gamma = 0.1$. In both networks, the local clustering coefficients of the nodes in the LCC of $\mathbf{A}^{(\mathrm{md})}$ are again markedly larger than the baseline values in the original networks in Fig. 1, although this is now less extreme than when $\gamma = 0.1$. In (a) and (b), "Comm." stands for "Community."

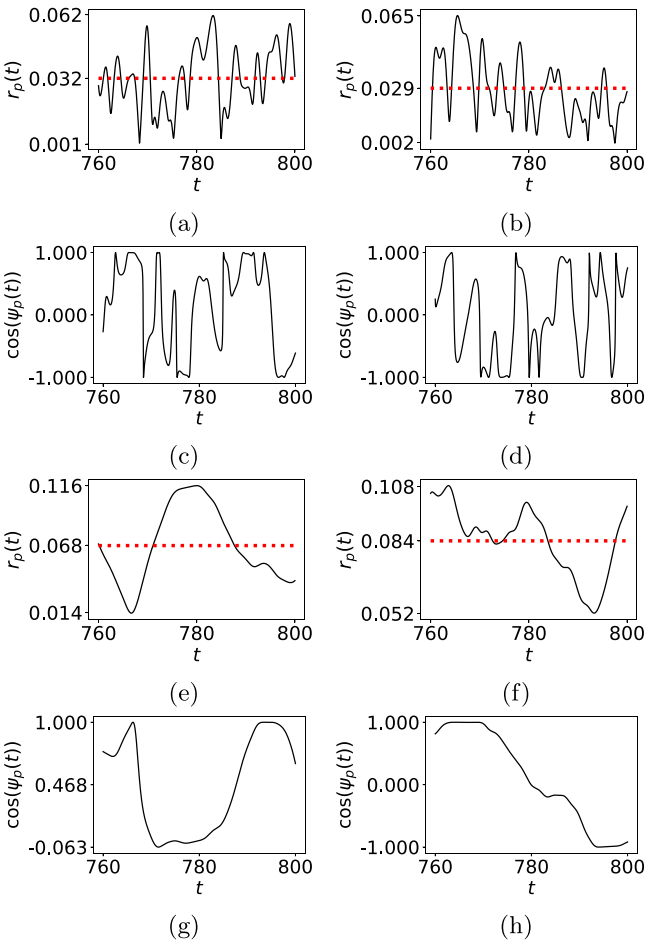

FIG. 7. The (a,b,e,f) order-parameter magnitude $r_p(t)$ and (c,d,g,h) real part $\cos[\psi_p(t)]$ of the order-parameter phase for the BA and WS networks when $\gamma = 1$. We show our results for (a,c,e,g) a BA network and (b,d,f,h) a WS network. Our results for the original two networks are in (a)–(d) and our results for the LCCs of $\mathbf{A}^{(\mathrm{md})}$ of these networks are in (e)–(h). The horizontal line in (a), (b), (e), and (f) indicates the mean $\langle r_p(\cdot)\rangle$ of the order-parameter magnitude. In both the BA network and the WS network, the LCC of $\mathbf{A}^{(\mathrm{md})}$ has much stronger frequency locking (as one can see by the slow variation in both $r_p(t)$ and $\cos[\psi_p(t)]$) than in the original networks.

yields a complicated DMD spectra, with approximately 340 eigenvalues for each network, in Figs. 8(c) and 8(d). The community counts in Figs. 9(a) and 9(b) reveal that even the LCCs of $\mathbf{A}^{(\mathrm{md})}$ have only seven oscillators for each network. However, as one can see in Figs. 9(c) and 9(d), these LCCs still have larger local clustering coefficients than the original networks.

In Fig. 10, we see that the functional community of oscillators that consists of the LCC of $\mathbf{A}^{(\mathrm{md})}$ is markedly closer to frequency locking than is the case for those oscillators in the original network. This arises from the larger, less oscillatory values of r_p and less erratic dynamics of $\cos[\psi_p(t)]$ in Figs. 10(e)–10(h) than in Figs. 10(a)–10(d). Therefore, although few oscillators are in these communities when $\gamma = 10$ (with the vast majority of oscillators wandering incoherently), these few oscillators are much closer to being locked than in the original network.

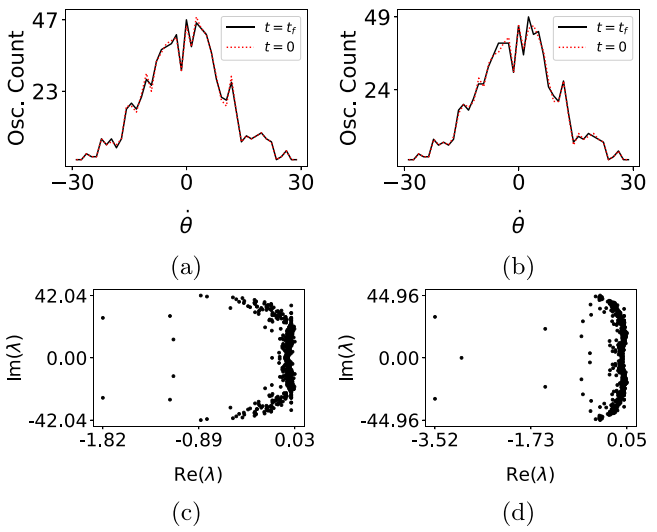

FIG. 8. The (a,b) frequency distributions and (c,d) DMD spectra for strongly heterogeneous Kuramoto oscillators (i.e., when $\gamma = 10$). We show our results for (a,c) a BA network and (b,d) a WS network. We plot the distributions of $\dot{\theta}_j$ at $t = 0$ and $t = t_f$ for (a) a BA network and (b) a WS network. It is now difficult to distinguish between the networks from the two models, although we still observe slight differences in the DMD spectra. In (a) and (b), "Osc." stands for "Oscillator."

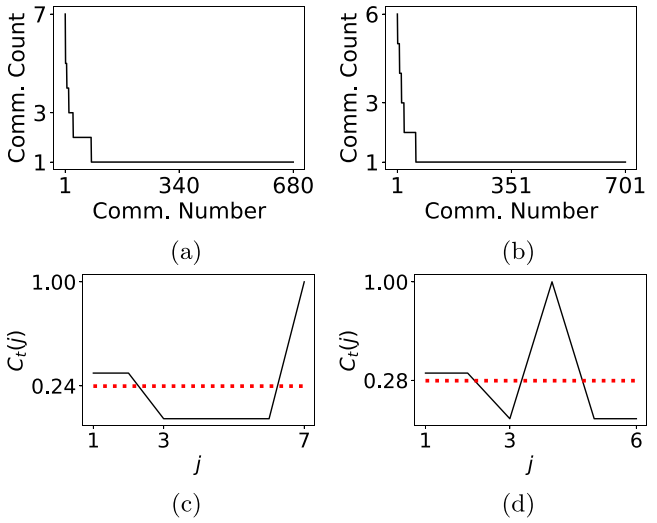

FIG. 9. The (a,b) numbers of nodes in each connected component of $\mathbf{A}^{(\mathrm{md})}$ and the (c,d) local clustering coefficients of the nodes in the LCC of $\mathbf{A}^{(\mathrm{md})}$ of the BA and WS networks for oscillators with $\gamma = 10$. The horizontal line in (c) and (d) indicates the mean value of the local clustering coefficient. We show our results for (a,c) a BA network and (b,d) a WS network. The large spread in oscillator frequencies for these strongly heterogeneous oscillators prevents the formation of large LCCs, and no community has more than a few oscillators in it. Nevertheless, the local clustering coefficients of the nodes in the LCC of $\mathbf{A}^{(\mathrm{md})}$ are still larger than the baseline values in the original networks in Fig. 1. In (a) and (b), "Comm." stands for "Community."

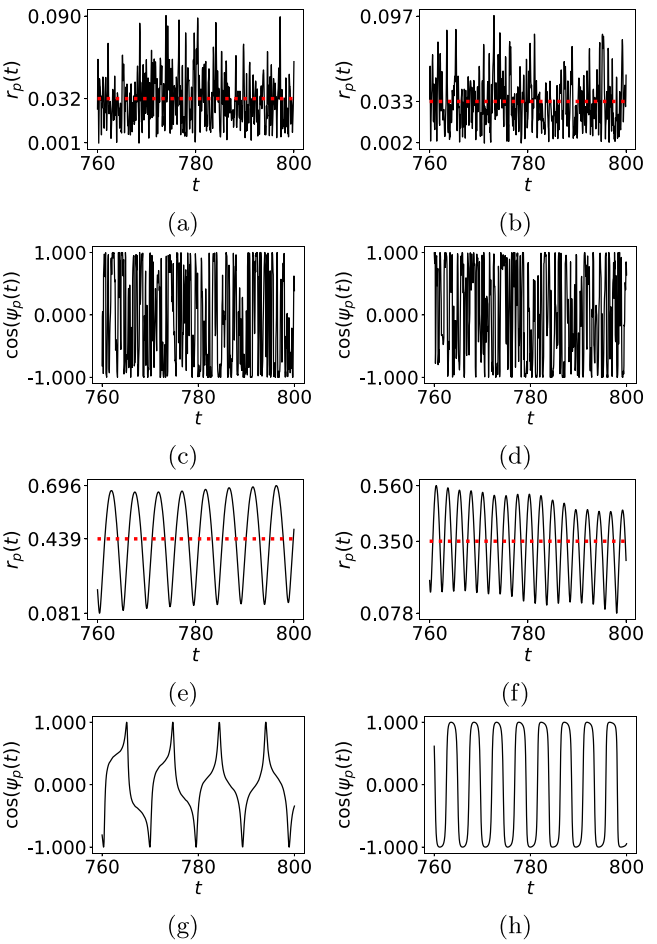

FIG. 10. The (a,b,e,f) order-parameter magnitude $r_p(t)$ and (c,d,g,h) real part $\cos[\psi_p(t)]$ of the order-parameter phase for the BA and WS networks when $\gamma = 10$. We show our results for (a,c,e,g) a BA network and (b,d,f,h) a WS network. Our results for the original two networks are in (a)–(d) and our results for the LCCs of $\mathbf{A}^{(\mathrm{md})}$ of these networks are in (e)–(h). The horizontal line in (a), (b), (e), and (f) indicates the mean $\langle r_p(\cdot)\rangle$ of the order-parameter magnitude. The networks from both models have more frequency locking in the LCC of $\mathbf{A}^{(\mathrm{md})}$ than in the original networks.

D. Community dynamics

Now that we have established our method of generating functional communities of oscillators in a manner that is consistent with their dynamics, we compare how communities form when $\gamma = 0.1$ (i.e., the case of weakly heterogeneous oscillators) in the two RGMs. In Fig. 11, we see that community formation occurs in starkly different ways in the two models, as the BA network has a far weaker tendency towards community coalescence than the WS network. However, in both of these networks, we observe coalescing communities that eventually cease to propagate. Our notion of functional communities allows us to examine their

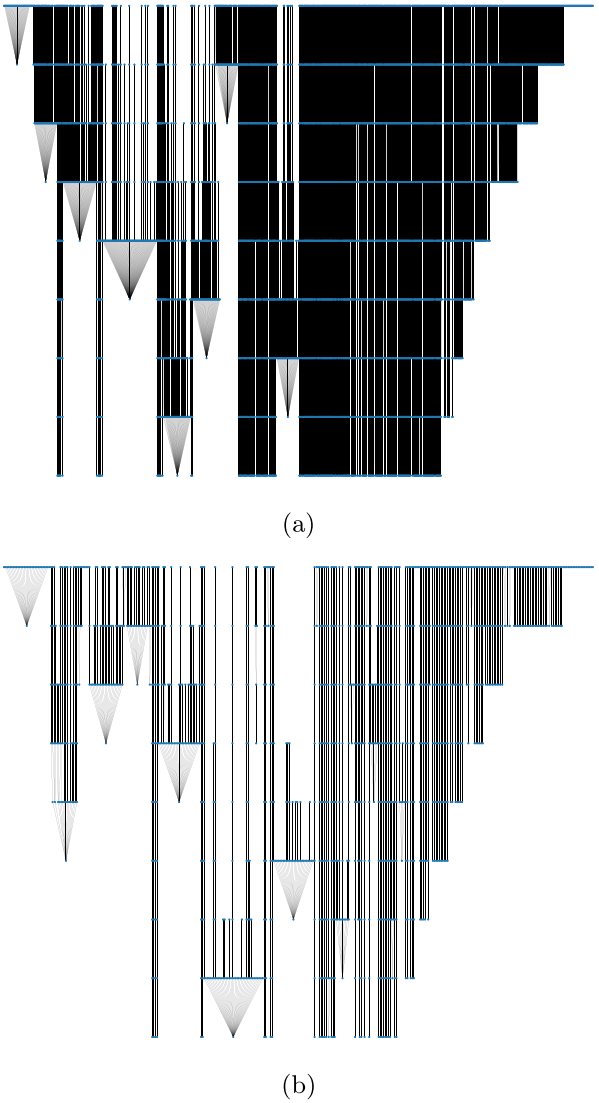

FIG. 11. Community-relationship graphs $\mathcal{G}_{t_i,t_f}$ for weakly heterogeneous oscillators (i.e., $\gamma = 0.1$) for (a) a BA network and (b) a WS network for the time interval [600, 800] (so $t_i = 600$ and $t_f = 800$) across overlapping intervals of 40 time units (i.e., with $n_w = 5$). Time runs from top to bottom. This plot illustrates how functional communities, which consist of oscillators with similar dynamics, evolve over time.

temporal evolution and thereby capture some details of the dynamics of the coupled oscillators.

V. CONCLUSIONS AND DISCUSSION

Using dynamic-mode decomposition, we developed a versatile approach for identifying functional communities of heterogeneous coupled oscillators on networks. We examined coupled Kuramoto oscillators on networks that we constructed from random-graph models, and we identified functional communities of oscillators based on how much they exhibit phase locking and frequency locking. These functional communities arise through clustering in graphs of community relationships, and the clustering coefficients in these graphs are larger than the clustering coefficients in the original networks. Unsurprisingly, in concert with associated synchronization properties, these functional communities are stronger for networks of weakly heterogeneous oscillators than they are for networks of strongly coupled oscillators. The community-relationship graphs, which take the form of forests, encode the interactions between the oscillators and provide a way to visualize their complicated dynamics over time. We observed coalescence of communities when oscillators exhibit phase locking or frequency locking. Additionally, from our community-relationship graphs, we observed that BA networks and WS networks yield functional communities with different community-coalescence properties.

Our results are promising, and it is important to test them in increasingly challenging scenarios. The Kuramoto model is very well-studied [6], so it useful to examine it as a starting point for developing approaches like ours. However, it is desirable to challenge our approach with other models of coupled oscillators, other dynamical systems, and time-series output of natural observations and laboratory experiments. Additionally, we only examined coupled oscillators with uncorrelated natural frequencies, and it will be fascinating to apply our approach to examine dynamics in the presence of such correlations.

ACKNOWLEDGMENT

We thank Arkady Pikovsky for helpful comments.

[1] M. E. J. Newman, *Networks*, 2nd ed. (Oxford University Press, Oxford, UK, 2018).

[2] M. A. Porter and J. P. Gleeson, *Dynamical Systems on Networks: A Tutorial*, Frontiers in Applied Dynamical Systems: Reviews and Tutorials, Vol. 4 (Springer International Publishing, Cham, Switzerland, 2016).

[3] I. Z. Kiss, J. C. Miller, and P. L. Simon, *Mathematics of Epidemics on Networks: From Exact to Approximate Models* (Springer International Publishing, Cham, Switzerland, 2017).

[4] S. Lehmann and Y.-Y. Ahn, *Complex Spreading Phenomena in Social Systems: Influence and Contagion in Real-World Social Networks* (Springer International Publishing, Cham, Switzerland, 2018).

[5] A. Arenas, A. Díaz-Guilera, J. Kurths, Y. Moreno, and C. Zhou, Synchronization in complex networks, Phys. Rep. **469**, 93 (2008).

[6] F. A. Rodrigues, T. K. D. M. Peron, P. Ji, and J. Kurths, The Kuramoto model in complex networks, Phys. Rep. **610**, 1 (2016).

[7] M. A. Porter, Nonlinearity + networks: A 2020 vision, in *Emerging Frontiers in Nonlinear Science*, edited by P. G. Kevrekidis, J. Cuevas-Maraver, and A. Saxena, Nonlinear Systems and Complexity, Vol. 32 (Springer International Publishing, Cham, Switzerland, 2020), pp. 131–159.

[8] M. A. Porter, J.-P. Onnela, and P. J. Mucha, Communities in networks, Not. Am. Math. Soc. **56**, 1082 (2009).

[9] S. Fortunato and D. Hric, Community detection in networks: A user guide, Phys. Rep. **659**, 1 (2016).
[10] A. Arenas, A. Díaz-Guilera, and C. J. Pérez-Vicente, Synchronization Reveals Topological Scales in Complex Networks, Phys. Rev. Lett. **96**, 114102 (2006).
[11] D. S. Bassett, M. A. Porter, N. F. Wymbs, S. T. Grafton, J. M. Carlson, and P. J. Mucha, Robust detection of dynamic community structure in networks, Chaos **23**, 013142 (2013).
[12] C. R. Shalizi, M. F. Camperi, and K. L. Klinkner, Discovering functional communities in dynamical networks, in *Statistical Network Analysis: Models, Issues, and New Directions*, edited by E. M. Airoldi, D. M. Blei, S. E. Fienberg, A. Goldenberg, E. P. Xing, and A. X. Zheng, Computer Communication Networks and Telecommunications Vol. 4503 (Springer-Verlag, Heidelberg, Germany, 2007), pp. 140–157.
[13] S. Chauhan, M. Girvan, and E. Ott, A network function-based definition of communities in complex networks, Chaos **22**, 033129 (2012).
[14] S. Boccaletti, M. Ivanchenko, V. Latora, A. Pluchino, and A. Rapisarda, Detecting complex network modularity by dynamical clustering, Phys. Rev. E **75**, 045102(R) (2007).
[15] J. Wu and Y. Jiao, Clustering dynamics of complex discrete-time networks and its application in community detection, Chaos **24**, 033104 (2014).
[16] P. Schmid, Dynamic mode decomposition of numerical and experimental data, J. Fluid Mech. **656**, 5 (2010).
[17] I. Mezić, Spectral properties of dynamical systems, model reduction, and decompositions, Nonlinear Dyn. **41**, 309 (2005).
[18] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems* (Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2016).
[19] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley, Modal analysis of fluid flows: An overview, AIAA J. **55**, 4013 (2017).
[20] S. Le Clainche and J. M. Vega, Analyzing nonlinear dynamics via data-driven dynamic mode decomposition-like methods, Complexity **2018**, 6920783 (2018).
[21] S. L. Brunton, M. Budišić, E. Kaiser, and J. N. Kutz, Modern Koopman theory for dynamical systems, arXiv:2102.12086.
[22] J. M. Kunert-Graf, K. M. Eschenburg, D. J. Galas, J. N. Kutz, S. D. Rane, and B. W. Brunton, Extracting reproducible time-resolved resting state networks, Front. Comput. Neurosci. **13**, 75 (2019).
[23] See https://github.com/cwcurtis/DMD_Community_Detect.
[24] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition, J. Nonlinear Sci. **25**, 1307 (2015).
[25] M. O. Williams, C. W. Rowley, and I. G. Kevrekidis, A kernel-based method for data-driven Koopman spectral analysis, J. Comput. Dyn. **2**, 247 (2015).
[26] C. D. Martin and M. A. Porter, The extraordinary SVD, Am. Math. Mon. **119**, 838 (2012).
[27] M. S. Hemati, C. W. Rowley, E. A. Deem, and L. N. Cattafesta, De-biasing the dynamic mode decomposition for applied Koopman spectral analysis of noisy data, Theor. Comput. Fluid Dyn. **31**, 349 (2017).
[28] H. Zhang, S. T. M. Dawson, C. W. Rowley, E. A. Deem, and L. N. Cattafesta, Evaluating the accuracy of the dynamic mode decomposition, J. Comput. Dyn. **7**, 35 (2020).
[29] A. V. Pimenova, D. S. Goldobin, M. Rosenblum, and A. Pikovsky, Interplay of coupling and common noise at the transition to synchrony in oscillator populations, Sci. Rep. **6**, 38518 (2016).
[30] C. M. Topaz, L. Ziegelmeier, and T. Halverson, Topological data analysis of biological aggregation models, PLoS One **10**, e0126383 (2015).
[31] J. Gómez-Gardeñes, Y. Moreno, and A. Arenas, Paths to Synchronization on Complex Networks, Phys. Rev. Lett. **98**, 034101 (2007).
[32] L. Buzna, S. Lozano, and A. Díaz-Guilera, Synchronization in symmetric bipolar population networks, Phys. Rev. E **80**, 066120 (2009).
[33] I. Sendiña-Nadal, J. M. Buldú, I. Leyva, and S. Boccaletti, Phase locking induces scale-free topologies in networks of coupled oscillators, PLoS One **3**, e2644 (2008).
[34] J. G. Restrepo, E. Ott, and B. R. Hunt, Onset of synchronization in large networks of coupled oscillators, Phys. Rev. E **71**, 036151 (2005).
[35] G. Palla, A.-L. Barabási, and T. Vicsek, Quantifying social group evolution, Nature (London) **446**, 664 (2007).
[36] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, Community structure in time-dependent, multiscale, and multiplex networks, Science **328**, 876 (2010).
[37] P. Holme, Modern temporal network theory: A colloquium, Eur. Phys. J. B **88**, 234 (2015).
[38] D. J. Fenn, M. A. Porter, P. J. Mucha, M. McDonald, S. Williams, N. F. Johnson, and N. S. Jones, Dynamical clustering of exchange rates, Quant. Financ. **12**, 1493 (2012).
[39] A.-L. Barabási and R. Albert, Emergence of scaling in random networks, Science **286**, 509 (1999).
[40] D. J. Watts and S. H. Strogatz, Collective dynamics of 'small-world' networks, Nature (London) **393**, 440 (1998).
[41] G. Haller, Lagrangian coherent structures, Annu. Rev. Fluid Mech. **47**, 137 (2015).
[42] M. R. Allshouse and T. Peacock, Lagrangian based methods for coherent structure detection, Chaos **25**, 097617 (2015).
[43] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2012).