



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Think locally, act locally: Detection of small, medium-sized, and large communities in large networks

Lucas G. S. Jeub, Prakash Balachandran, Mason A. Porter, Peter J. Mucha, and Michael W. Mahoney

Phys. Rev. E **91**, 012821 — Published 26 January 2015

DOI: [10.1103/PhysRevE.91.012821](https://doi.org/10.1103/PhysRevE.91.012821)

Think Locally, Act Locally: The Detection of Small, Medium-Sized, and Large Communities in Large Networks

Lucas G. S. Jeub,¹ Prakash Balachandran,² Mason A. Porter,^{1,3} Peter J. Mucha,^{4,5} and Michael W. Mahoney^{6,7}

¹*Oxford Centre for Industrial and Applied Mathematics,
Mathematical Institute, University of Oxford, OX2 6GG, UK*

²*Department of Mathematics & Statistics, Boston University, Boston, MA 02215, USA*

³*CABDyN Complexity Centre, University of Oxford, Oxford, OX1 1HP, UK*

⁴*Department of Applied Physical Sciences, University of North Carolina, Chapel Hill, NC 27599-3216, USA*

⁵*Institute for Advanced Materials, Nanoscience & Technology,
University of North Carolina, Chapel Hill, NC 27599-3216, USA*

⁶*International Computer Science Institute, Berkeley, CA 94704*

⁷*Department of Statistics, University of California at Berkeley, Berkeley, CA 94720*

It is common in the study of networks to investigate intermediate-sized (or “meso-scale”) features to try to gain an understanding of network structure and function. For example, numerous algorithms have been developed to try to identify “communities,” which are typically construed as sets of nodes with denser connections internally than with the remainder of a network. In this paper, we adopt a complementary perspective that “communities” are associated with bottlenecks of locally-biased dynamical processes that begin at seed sets of nodes, and we employ several different community-identification procedures (using diffusion-based and geodesic-based dynamics) to investigate community quality as a function of community size. Using several empirical and synthetic networks, we identify several distinct scenarios for “size-resolved community structure” that can arise in real (and realistic) networks: (i) the best small groups of nodes can be better than the best large groups (for a given formulation of the idea of a good community); (ii) the best small groups can have a quality that is comparable to the best medium-sized and large groups; and (iii) the best small groups of nodes can be worse than the best large groups. As we discuss in detail, which of these three cases holds for a given network can make an enormous difference when investigating and making claims about network community structure, and it is important to take this into account to obtain reliable downstream conclusions. Depending on which scenario holds, one may or may not be able to successfully identify “good” communities in a given network (and good communities might not even exist for a given community quality measure), the manner in which different small communities fit together to form meso-scale network structures can be very different, and processes such as viral propagation and information diffusion can exhibit very different dynamics. In addition, our results suggest that, for many large realistic networks, the output of locally-biased methods that focus on communities that are centered around a given seed node might have better conceptual grounding and greater practical utility than the output of global community-detection methods. They also illustrate subtler structural properties that are important to consider in the development of better benchmark networks to test methods for community detection.

PACS numbers: 89.75.Fb, 89.75.Hc, 05.10.-a

I. INTRODUCTION

Many physical, technological, biological, and social systems can be modeled as networks, which in their simplest form are represented by graphs. A (static and single-layer) graph consists of a set of entities (called “vertices” or “nodes”) and pairwise interactions (called “edges” or “links”) between those vertices [1–3]. Graphical representations of data have led to numerous insights in the natural, social, and information sciences; and the study of networks has in turn borrowed ideas from all of these areas [4].

In general, networks can be described using a combination of local, global, and “meso-scale” perspectives. To investigate meso-scale structures—i.e., intermediate-sized structures that are responsible for “coupling” local properties, such as whether triangles close, and global properties such as graph diameter—a fundamental prim-

itive in many applications entails partitioning graphs into meaningful and/or useful sets of nodes [3]. The most popular form of such a partitioning procedure, in which one attempts to find relatively dense sets of nodes that are relatively sparsely connected to other sets, is known as “community detection” [5–7]. Myriad methods have been developed to algorithmically detect communities [5, 6]; and these efforts have led to insights in applications such as committee and voting networks in political science [8–10], friendship networks at universities and other schools [11–13], protein-protein interaction networks [14], granular materials [15], amorphous materials [16], brain and behavioral networks in neuroscience [17–19], collaboration patterns [20], human communication networks [21, 22], human mobility patterns [23], and so on.

The motivation for the present work is the observation that it can be very challenging to find meaningful

medium-sized or large communities in large networks [24–26]. Much of the large body of work on algorithmically identifying communities in networks has been applied successfully either to find communities in small networks or to find small communities in large networks [5, 6, 25], but it has been much less successful at finding meaningful medium-sized and large communities in large networks [27]. There are many reasons that it is difficult to find “good” large communities in large networks. We discuss several such reasons in the following paragraphs.

First, although it is typical to think about communities as sets of nodes with “denser” interactions among its members than between its members and the rest of a network, the literature contains neither a consensus definition of community nor a consensus on a precise formalization of what constitutes a “good” community [5, 6].

Second, the popular formalizations of a “community” are computationally intractable, and there is little precise understanding or theoretical control on how closely popular heuristics to compute communities approximate the exact answers in those formulations [28, 29]. Indeed, community structure itself is typically “defined” operationally via the output of a community-detection algorithm, rather than as the solution to a precise optimization problem or via some other mathematically precise notion [5, 6].

Third, many large networks are extremely sparse [25] and thus have complicated structures that pose significant challenges for the algorithmic detection of communities via the optimization of objective functions [29]. This is especially true when attempting to develop algorithms that scale well enough to be usable in practice on large networks [6, 25, 30].

Fourth, the fact that it is difficult to visualize large networks complicates the validation of community-detection methods in such networks. One possible means of validation is to compare algorithmically-obtained communities with known “ground truth” communities. However, notions of ground truth can be weak in large networks [12, 25, 31], and one rarely possesses even a weak notion of ground truth for most networks. Indeed, in many cases, one should not expect a real (or realistic) large network to possess a single feature that (to leading order) dominates large-scale latent structure in a network. Thus, comparing the output of community-detection algorithms to “ground truth” in practice is most appropriate for obtaining coarse insights into how a network might be organized into social or functional groups of nodes [11]. Alternatively, different notions and/or formalizations of “community” concepts might be appropriate in different contexts [5, 6, 32–34], so it is desirable to formulate flexible methods that can incorporate different perspectives.

Fifth, community-detection algorithms often have subtle and counterintuitive properties as a function of sizes of their inputs and/or outputs. For example, the community-size “resolution limit” of the popular modularity objective function is a fundamental consequence

of the additive form of that objective function, but it only became obvious to people after it was explicitly pointed out [35].

Motivated by these observations, we consider the question of community quality as a function of the size (i.e., number of nodes) of a purported community. That is, we are concerned with questions such as the following. (1) What is the relationship between communities of different sizes in a given network? In particular, for a given network and a given community-quality objective, are larger communities “better” or “worse” than smaller communities? (2) What is an appropriate way to think about medium-sized and large communities in large networks? In particular, how do smaller communities “fit together” into medium-sized and larger communities? (3) More generally, what effect do the answers to these questions have on downstream tasks that are of primary concern when modeling data using networks? For example, what effect do they have on processes such as viral propagation or the diffusion of information on networks?

By considering a suite of networks and using several related notions of community quality, we identify several scenarios that can arise in realistic networks.

1. **Small communities are better than large communities.** In this first scenario, for which there is an upward-sloping *network community profile* (NCP; see the discussion below), a network has small groups of nodes that correspond more closely than any large groups to intuitive ideas of what constitutes a good community.
2. **Small and large communities are similarly good or bad.** In this second scenario, for which an NCP is roughly flat, the most community-like small groups of nodes in a network have similar community quality scores to the most community-like large groups.
3. **Large communities are better than small communities.** In this third scenario, for which an NCP is downward-sloping, a network has large groups of nodes that are more community-like (i.e., “better” in some sense) than any small groups.

Although the third scenario is the one that has an intuitive isoperimetric interpretation and thus corresponds most closely with peoples’ intuition when they develop and validate community-detection algorithms, one of our main conclusions is that most large realistic networks correspond to the first or second scenarios. This is consistent with recent results on network community structure using related approaches [24–26] as well as somewhat different approaches [31, 36], and it also helps illustrate the importance of considering community structures with groups that have large overlaps. For more on this, see our discussions below.

One of the main tools that we use to justify the above observations and to interpret the implications of community structure in a network is a *network community pro-*

file (NCP), which was originally introduced in Ref. [25]. Given a community “quality” score—i.e., a formalization of the idea of a “good” community—an NCP plots the score of the best community of a given size as a function of community size. The authors of Ref. [24, 25] considered the community quality notion of conductance and employed various algorithms to approximate it. In subsequent work [26], many other notions of community quality have also been used to compute NCPs.

In the present paper, we compute NCPs using three different procedures to identify communities.

1. **Diffusion-based dynamics.** First, we consider a diffusion-based dynamics (called the ACLCUT method; see the discussion below) from the original NCP analysis [25] that has an interpretation that good communities correspond to bottlenecks in the associated dynamics.
2. **Spectral-based optimization.** Second, we consider a spectral-based optimization rule (called the MOV CUT method; see below) that is a locally-biased analog of the usual global spectral graph partitioning problem [37].
3. **Geodesic-based dynamics.** Finally, we consider a geodesic-based spreading process (called the EGONET method; see the discussion below) that has an interpretation that nodes in a good community are connected by short paths that emanate from a seed node [38].

We describe these three procedures in more detail in Appendix B. For now, we note that the first and the third procedures have a natural interpretation as defining communities operationally as the output of an underlying dynamics, and the first and second procedures allow us to compare this operational approach with an optimization-based approach.

Viewed from this perspective, the computation of network community structure depends fundamentally on three things: actual network structure, the dynamics or application of interest, and the initial conditions or network region of interest. Although there are differences between the aforementioned three community-identification methods, these methods all take the perspective that a network’s community structure depends not only on the connectivity of its nodes but also on (1) the region of a large network in which one is interested and (2) the application of interest. The perspective in point (1) contrasts with the prevalent view of community structure as arising simply from network structure [5, 6], but it is consistent with the notion of dynamical systems depending fundamentally on their initial conditions, and it is crucial in many applications (e.g., both social [39, 40] and biological contagions [41–43]).

For example, Facebook’s Data Team and its collaborators have demonstrated that one can view Facebook as a collection of egocentric networks that have been patched together into a network whose global structure

is very sparse [44, 45]. The above three community-identification methods have the virtue of combining the prevalent structural perspective with the idea that one is often interested in structure that is located “near” (in terms of both network topology and edge weights) an exogenously-specified “seed set” of nodes [46]. The perspective in point (2) underscores the fact that one should not expect answers to be “universal.” The differences between the aforementioned three methods lie in the specific dynamical processes that underlie them. We also note that, although we focus on the measure of community quality known as “conductance” (which is intimately related to the problem of characterizing the mixing rates of random walks [47]), one can view other quality functions (e.g., based on non-conservative dynamics [48–50] or geodesic-based dynamics [38]) as solving other problems, and they thus can reveal different aspects of community structure in networks.

The global NCPs that we compute from the three community-identification procedures are rather similar in some respects, suggesting that the characteristic features of NCPs are actual features of networks and not just artifacts of a particular way of sampling local communities. However, we observe significant differences in their local behaviors because they are based on different dynamical processes. In concert with other recent work (e.g., [34, 51, 52]), our results with these three procedures suggest that “local” methods that focus on finding communities centered around an exogenously-specified seed node (or seed set of nodes) might have better theoretical grounding and more practical utility than other methods for community detection.

Our “local” (and “size-resolved”) perspective on community structure also yields several other interesting insights. By design, it allows us to discern how community structure depends both on the seed node and on the size scales and time scales of a dynamical process running on a network. Similar perspectives were discussed in recent work on detecting communities in networks using Markov processes [9, 32–34, 53–55], and our approach is in the spirit of research on dynamical systems more generally, as bottlenecks to diffusion and other dynamics depend fundamentally on initial conditions. Local information algorithms are also an important approach for many other optimization problems and for practical purposes such as friend recommendation systems in online social networks [56]. Moreover, taking a local perspective on community structure is also consistent with the sociological idea of egocentric networks (and with real-world experience of individuals, such as users of Facebook [44, 45], who experience their personal neighborhood of a social network). The local community experienced by a given node should be similarly locally-biased, and we demonstrate this feature quantitatively for several real networks. Using our perspective, we also demonstrate subtle yet fundamental differences between different networks: some networks have high-quality communities of different sizes (especially small ones), whereas others do not possess commu-

nities of any size that give bottlenecks to diffusion-based dynamics. This is consistent with, and helps explain, prior direct observations of networks in which algorithmically computed communities seemed to have little or no effect on several dynamical processes [57]. More generally and importantly, whether small or large communities are “better” with respect to some measure of community quality has significant consequences not only for algorithms that attempt to identify communities but also for the dynamics of processes such as viral propagation and information diffusion.

The rest of this paper is organized as follows. Because our approach to examining network communities is uncommon in the physics literature, we start in Section II with an informal description of our approach. We then introduce NCPs in Section III. In Section IV, we present our main empirical results on community quality as a function of size, and we provide a detailed comparison of our three community-identification procedures when applied to real networks. This illustrates the three distinct scenarios of community quality versus community size that we described above. In Section V, we illustrate the behavior of these methods on the well-known LFR benchmark networks that are commonly used to evaluate the performance of community-detection techniques. We find that their NCPs have a characteristic shape for a wide range of parameter values and are unable to reproduce the different scenarios that one observes for real networks. We then conclude in Section VI with a discussion of our results. In Appendix A, we provide a brief discussion of expander graphs (a.k.a. “expanders”). In Appendix B, we describe the three specific procedures that we use to identify communities in detail. Appendices C and D contain empirical results for the two methods that we mentioned but did not discuss in detail in Section IV.

II. BACKGROUND AND PRELIMINARIES

In this section, we describe some background and preliminaries that provide the framework that we use to interpret our results on size-resolved community structure in Sections IV and V. We start in Section II A by defining the notation that we use throughout this paper, and we continue in Section II B with a brief discussion of possible ways that a network might “look like” if one is interested in its meso-scale or large-scale structure. To convey the basic idea of our approach, much of our discussion in this section is informal. In later sections, we will make these ideas more precise.

A. Definitions and Notation

We represent each of the networks that we study as an undirected graph. We consider both weighted and unweighted graphs.

Let $G = (V, E, w)$ be a connected and undirected graph with node set V , edge set E , and set w of weights on the edges. Let $n = |V|$ denote the number of nodes, and let $m = |E|$ denote the number of edges. The edge $\{i, j\}$ has weight w_{ij} . Let $A = A_G \in \mathbb{R}^{n \times n}$ denote the (weighted) adjacency matrix of G . Its components are $A_G(i, j) = w_{ij}$ if $\{i, j\} \in E$ and $A_G(i, j) = 0$ otherwise. The matrix $D = D_G \in \mathbb{R}^{n \times n}$ denotes the diagonal degree matrix of G . Its components are $D_G(i, i) = d_i = \sum_{\{i, j\} \in E} w_{ij}$, where d_i is called the “strength” or “weighted degree” of node i . The combinatorial Laplacian of G is $L_G = D_G - A_G$, and the normalized Laplacian of G is $\mathcal{L}_G = D_G^{-1/2} L_G D_G^{-1/2}$.

A path P in G is a sequence of edges $P = \{\{i_k, j_k\}\}_{k=1}^s$, such that $j_k = i_{k+1}$ for $k = 1, \dots, s-1$. The length of path P is $|P| = \sum_{\{i, j\} \in P} l_{ij}$, where l_{ij} is the length of the edge that connects nodes i and j . For an unweighted network, $l_{ij} = 1$ for all edges. For weighted networks, w_{ij} is a measure of closeness of the tie between nodes i and j , a common choice for l is $l_{ij} = \frac{1}{w_{ij}}$. Let \mathcal{P}_{ij} be the set of all paths between i and j . The geodesic distance $\Delta_{ij} = \min_{P \in \mathcal{P}_{ij}} |P|$ between nodes i and j is the length of a shortest path between i and j . The k -neighborhood $N_k(i) = \{j \in V : \Delta_{ij} \leq k\}$ of i is the set of all nodes that are at most a distance k away from i , and the k -neighborhood of a set of nodes S is $N_k(S) = \bigcup_{i \in S} N_k(i)$.

B. What Can Networks “Look Like”?

Before examining real networks, we start with the following question: What are possible ways that a network can “look like,” very roughly if one “squints” at it? This question is admittedly vague, but the answer to it governs how small-scale network structure “interacts” with large-scale network structure, and it informs researchers’ intuitions and the design decisions that they make when analyzing networks (and when developing methods to analyze networks). As an example of this idea, it should be intuitively clear that if one “squints” at the nearest-neighbor \mathbb{Z}^2 network (i.e., the uniform lattice of pairs of integers on the Euclidean plane), then they “look like” the Euclidean plane \mathbb{R}^2 . Distances are approximately preserved, and up to boundary conditions and discretization effects, dynamical processes on one approximate the analogous dynamic processes on the other. In the fields of geometric group theory and coarse geometry, this intuitive connection between \mathbb{Z}^2 and \mathbb{R}^2 has been made precise using the notions of coarse embeddings and quasi-isometries [58].

Establishing quasi-isometric relationships on networks that are expander graphs (a.k.a. “expanders”; see Appendix A) is technically brittle [60]. Thus, for the present informal discussion, we rely on a simpler notion. Suppose that we are interested in the “best fit” of the adjacency




FIG. 1. Idealized block models of network adjacency matrices; darker blocks correspond to denser connections among its component nodes. Figure 1a illustrates a low-dimensional “hot dog” or “pancake” structure; Fig. 1b illustrates a “core-periphery” structure; Fig. 1c illustrates an unstructured expander or complete graph; and Fig. 1d illustrates a bipartite graph. Our example networks are the Zachary Karate Club [59] in Fig. 1a and a realization of a random-graph block model in Figs. 1b–1d. For Fig. 1b we only show the largest connected component (LCC), whereas the networks in Figs. 1c and 1d are connected. The parameters for the block models are as follows: (b) $\alpha_{11} = 0.3$, $\alpha_{22} = 0.001$, $\alpha_{12} = 0.005$, $n_1 = 50$ nodes, and $n_2 = 950$ nodes (the LCC has 615 nodes); (c) $\alpha_{11} = \alpha_{22} = \alpha_{12} = 0.01$, and $n_1 + n_2 = 1000$ nodes; (d) $\alpha_{11} = \alpha_{22} = 0$, $\alpha_{12} = 0.02$, and $n_1 = n_2 = 500$ nodes.

matrix A to a 2×2 block matrix:

$$A = \begin{pmatrix} A_{11} & A_{12}^T \\ A_{12} & A_{22} \end{pmatrix},$$

where $A_{ij} = \alpha_{ij} \vec{1} \vec{1}^T$, where the “1-vector” $\vec{1}$ is a column vector of the appropriate dimension that contains a 1 in every entry and $\alpha_{ij} \in \mathbb{R}^+$. Thus, each block in A has uniform values for all its elements, and larger values of α_{ij} correspond to stronger interactions between nodes. The structure of A is then determined based on the relative sizes of α_{11} , α_{12} , and α_{22} . The various relative sizes of these three scalars have a strong bearing on the structure of the network associated with A . We illustrate several examples in Fig. 1. For the block models that we use for three of its panels, one block has n_1 nodes and the second block has n_2 nodes, and a node in block i is connected to a node in block j with probability α_{ij} [61].

- **Low-dimensional structure.** In Fig. 1a, we illustrate the case in which $\alpha_{11} \approx \alpha_{22} \gg \alpha_{12}$. In this case, each half of the network interacts with itself more densely than it interacts with the other half of the network. This “hot dog” or “pancake” structure corresponds to the situation in which there are two (or any number, in the case of networks more generally) dense communities of nodes that are reasonably well-balanced in the sense that each

community has roughly the same number of nodes. In this case, the network embeds relatively well in a one-dimensional, two-dimensional, or other low-dimensional space. Spectral clustering or other clustering methods often find meaningful communities in such networks, and one can often readily construct meaningful and interpretable visualizations of network structure.

- **Core-periphery structure.** In Fig. 1b, we illustrate the case in which $\alpha_{11} \gg \alpha_{12} \gg \alpha_{22}$. This is an example of a network with a density-based “core-periphery” structure [24, 25, 62–64]. In these cases, there is a core set of nodes that are relatively well-connected amongst themselves as well as to a peripheral set of nodes that interact very little amongst themselves.
- **Expander or complete graph.** In Fig. 1c, we illustrate the case in which $\alpha_{11} \approx \alpha_{12} \approx \alpha_{22}$. This corresponds to a network with little or no discernible structure. For example, if $\alpha_{11} = \alpha_{12} = \alpha_{22} = 1$, then the graph is a clique (i.e., the complete graph). Alternatively, if the graph is a constant-degree expander, then $\alpha_{11} \approx \alpha_{12} \approx \alpha_{22} \ll 1$. As discussed in Appendix A, constant-degree expanders are the metric spaces that embed least well in low-dimensional Euclidean spaces. In terms of the idealized block model in Fig. 1, they “look like” complete graphs, and partitioning them would not yield network structure that one should expect to construe as meaningful. Informally, they are largely unstructured when viewed at large size scales.
- **Bipartite structure.** In Fig. 1d, we illustrate the case in which $\alpha_{12} \gg \alpha_{11} \approx \alpha_{22}$. This corresponds to a bipartite or nearly-bipartite graph. Such networks arise, e.g., when there are two different types of nodes, such that one type of node connects only to (or predominantly to) nodes of the other type [65].

Most methods for algorithmic detection of communities have been developed and validated using the intuition that networks have some sort of low-dimensional structure [5, 25, 36]. As an example, consider the infamous Zachary Karate Club network [59], which we show in Fig. 1a. This well-known benchmark graph, which seems to be an almost obligatory example to discuss in papers that discuss community structure [66, 67], clearly “looks like” it has a nice low-dimensional structure. For example, there is a clearly identifiable left half and right half, and two-dimensional visualizations of the network (such as that in Fig. 1a) highlight that bipartition. Indeed, the Zachary Karate Club network possesses well-balanced and (quoting Herbert Simon [68]) “nearly decomposable” communities; and the nodes in each community are more densely connected to nodes in the same community than they are to nodes in the other community. Relatedly,

reordering the nodes of the Zachary Karate Club appropriately yields an adjacency-matrix representation with an almost block-diagonal structure with two blocks (as typified by the cartoon in Fig. 1a); and any reasonable community-detection algorithm should be able to find (exactly or approximately) the two communities.

Another well-known network that (slightly less obviously) “looks like” it has a low-dimensional structure is a so-called caveman network, which we illustrate later (in Fig. 2c). Arguably, a caveman network has many more communities than the Zachary Karate Club, so details such as whether an algorithm “should” split it into two or a somewhat larger number of reasonably well-balanced communities might be different than in the Zachary Karate Club network. However, a caveman network also has a natural well-balanced partition that respects intuitive community structure. Reasonable two-dimensional visualizations of this network (such as the one that we present in Fig. 2c) shed light on that structure; and any reasonable community-detection algorithm can be adjusted to find (exactly or approximately) the expected communities. In this paper, we will demonstrate that most realistic networks do *not* “look like” these small examples. Instead, realistic networks are often poorly-approximated by low-dimensional structures (e.g., with a small number of relatively well-balanced communities, each of which is more densely connected internally than it is with the rest of the network). Realistic networks often include substructures that more closely resemble core-periphery graphs or expander graphs (see Fig. 1b and Fig. 1c); and networks that partition into nice nearly-decomposable communities tend to be the exception rather than typical [24, 25, 36].

III. NETWORK COMMUNITY PROFILES (NCPS) AND THEIR INTERPRETATION

Recall from Section I that an NCP measures the quality of the best possible community of a given size as a function of the size of the purported community [24–26]. In this section, we provide a brief description of NCPS and how we will use it.

A. The Basic NCP: Measuring Size-Resolved Community Quality

We start with the definition of conductance and the original conductance-based definition of an NCP from Ref. [25], and we then discuss our extensions of such ideas. For more details on conductance and NCPS, see Refs. [25, 37, 69, 70]. If $G = (V, E, w)$ is a graph with weighted adjacency matrix A , then the “volume” between two sets S_1 and S_2 of nodes (i.e., $S_i \subset V$) equals the total weight of edges with one end in S_1 and one end

in S_2 . That is,

$$\text{vol}(S_1, S_2) = \sum_{i \in S_1} \sum_{j \in S_2} A_{ij}. \quad (1)$$

In this case, the “volume” of a set $S \subset V$ of nodes is

$$\text{vol}(S) = \text{vol}(S, V) = \sum_{i \in S} \sum_{j \in V} A_{ij}. \quad (2)$$

In other words, the set volume equals the total weight of edges that are attached to nodes in the set. The volume $\text{vol}(S, \bar{S})$ between a set S and its complement \bar{S} has a natural interpretation as the “surface area” of the “boundary” between S and \bar{S} . In this study, a set S is a hypothesized community. Informally, the conductance of a set S of nodes is the “surface area” of that hypothesized community divided by “volume” (i.e., size) of that community. From this perspective, studying community structure amounts to an exploration of the isoperimetric structure of G .

Somewhat more formally, the *conductance of a set of nodes* $S \subset V$ is

$$\phi(S) = \frac{\text{vol}(S, \bar{S})}{\min(\text{vol}(S), \text{vol}(\bar{S}))}. \quad (3)$$

Thus, smaller values of conductance correspond to better communities. The *conductance of a graph* G is the minimum conductance of any subset of nodes:

$$\phi(G) = \min_{S \subset V} \phi(S). \quad (4)$$

Computing the conductance $\phi(G)$ of an arbitrary graph is an intractable problem (in the sense that the associated decision problem is NP-hard [71]), but this quantity can be approximated by the second smallest eigenvalue λ_2 of the normalized Laplacian [69, 70].

If the “surface area to volume” (i.e., isoperimetric) interpretation captures the notion of a good community as a set of nodes that is connected more densely internally than with the remainder of a network, then computing the solution to Eq. (4) leads to the “best” (in this sense) community of any size in the network.

Instead of defining a community quality score in terms of the best community of any size, it is useful to define a community quality score in terms of the best community of a given size k as a function of the size k . To do this, Ref. [25] introduced the idea of a *network community profile (NCP)* as the lower envelope of the conductance values of communities of a given size:

$$\phi_k(G) = \min_{S \subset V, |S|=k} \phi(S). \quad (5)$$

An NCP plots a community quality score (which, as in Ref. [25], we take to be the set conductance of communities) of the best possible community of size k as a function of k . Clearly, it is also intractable to compute the quantity $\phi_k(G)$ in Eq. (5) exactly. Previous work has used

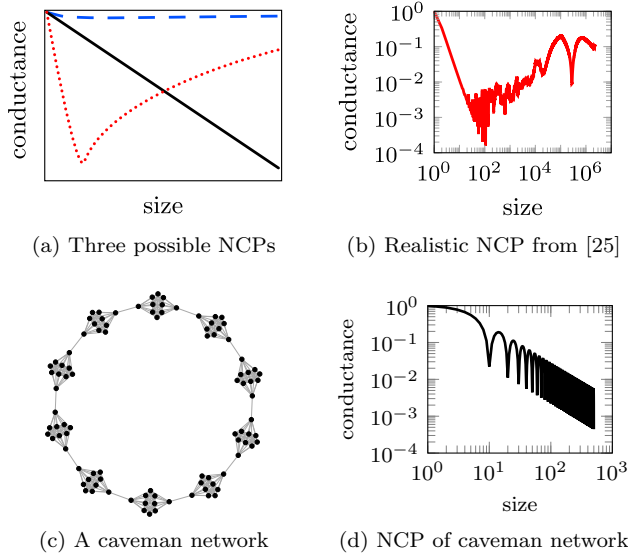


FIG. 2. Illustration of network community profiles (NCPs). (a) Stylized versions of possible shapes for an NCP: downward-sloping (black, solid), upward-sloping (red, dotted), and flat (blue, dashed). (b) NCP of a LIVEJOURNAL network that illustrates the characteristic upward-sloping NCP that is typical for many large empirical social and information networks [25]. (c) A toy “caveman network” with 10 cliques of 10 nodes each, where one edge from each clique has been rewired to create a ring [72]. (d) NCP for a similar caveman network with 100 cliques of 10 nodes each (the NCP for the network in panel (c) is identical for communities with fewer than 50 nodes), illustrating the characteristic downward-sloping NCP that is typical of networks that are embedded in a low-dimensional space.

spectral-based and flow-based approximation algorithms to approximate it [24–26].

To gain insight into how to understand an NCP and what it reveals about network structure, consider Fig. 2. In Fig. 2a, we illustrate three possible ways that an NCP can behave. In each case, we are using conductance as a measure of community quality.

- **Upward-sloping NCP.** In this case, small communities are “better” than large communities.
- **Flat NCP.** In this case, community quality is independent of size. (As illustrated in this figure, the quality tends to be comparably poor for all sizes.)
- **Downward-sloping NCP.** In this case, large communities are “better” than small communities.

For ease of visualization and computational considerations, we only show NCPs for communities up to half of the size of a network. An NCP for very large communities that we do not show in figures as a result of this choice roughly mirrors that for small communities, as the complement of a good small community is a good large

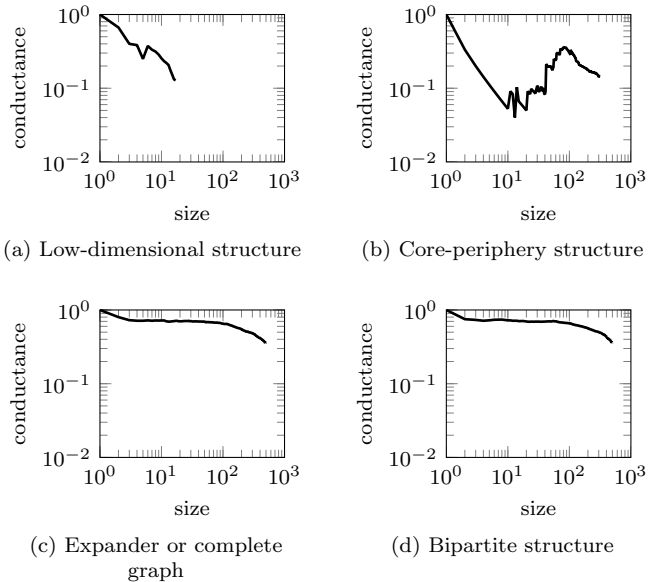


FIG. 3. Network community profiles (NCPs) of the idealized example networks from Fig. 1. (a) NCP for the Zachary Karate Club network. (b) NCP for an example network generated from a block model with core-periphery structure. (c) NCP for an Erdős-Rényi graph. (d) NCP for an example network generated from a bipartite block model.

community because of the inherent symmetry in conductance (see Eq. (3)).

In Fig. 2b, we show an NCP of a LiveJournal network from Ref. [25]. It demonstrates an empirical fact about a wide range of large social and information networks: there exist good small conductance-based communities, but there do not exist any good large conductance-based communities in many such networks. See Refs. [24–26, 37, 69, 70]) for more empirical evidence that large social and information networks tend not to have large communities with low conductances. On the contrary, Fig. 2c illustrates a small toy network—a so-called “caveman network”—formed from several small cliques connected by rewiring one edge from each clique to create a ring [72]. As illustrated by its downward-sloping NCP in Fig. 2d, this network possesses good conductance-based communities, and large communities are better than small ones. One obtains a similar downward-sloping NCP for the Zachary Karate Club network [59] as well as for many other networks for which there exist meaningful visualizations [25]. The wide use of networks that have interpretable visualizations (such as the Zachary Karate Club and planted partition models [73] with balanced communities) to help develop and evaluate methods for community detection and other procedures can lead to a strong selection bias when evaluating the quality of those methods.

We now consider the relationship between the phenomena illustrated in Fig. 2 and the idealized block models of Fig. 1. As a concrete example, Fig. 3 shows the NCPs

for the example networks in the right panels of Fig. 2.

First, note that the best partitions consist roughly of well-balanced communities in the low-dimensional case of Figs. 1a and 3a, and the “lowest” point on an NCP tends to be for large community sizes. Thus, an NCP tends to be downward-sloping.

Networks with pronounced core-periphery structure—i.e., networks that “look like” the example network in Fig. 1b—tend to have many good small communities but no equally good or better large communities. This situation arises in many large, extremely sparse networks [24–26]. The good small communities in such networks are sets of connected nodes in the extremely sparse periphery, and they do not combine to form good, large communities, as they are only connected via a set of core nodes with denser connections than the periphery. Thus, an NCP of a network with core-periphery structure tends to be upward-sloping, as illustrated in Figs. 1b and 3b. However, this observation does not apply to all networks with well-defined density-based core-periphery structure. If the periphery is sufficiently well-connected (though still much sparser than the core), then one no longer observes good, small communities. Such networks act like expanders from the perspective of the behavior of random walkers, so they have a flat NCP. One can generate examples of such networks by modifying the parameters of the block-model that we used to generate the example network in Fig. 1b [61].

For a complete graph or a degree-homogeneous expander (see Figs. 1c and 3c), all communities tend to have poor quality, so an NCP is roughly flat. (See Appendix A for a discussion of expander graphs.)

Finally, bipartite structure itself does not have any characteristic influence on an NCP. Instead, an NCP of a bipartite network reveals other structure present in a network. For the example network in Fig. 1d, the two types of nodes are connected uniformly at random, so its NCP (Fig. 3d) has the characteristic flat shape of an expander.

B. Robustness and Information Content of NCPs

It is important to discuss the robustness properties of NCPs. These are not obvious a priori, as the NCP is an extremal diagnostic. Importantly, though, the qualitative property of being downward-sloping, upward-sloping, or roughly flat is very robust to the removal of nodes and edges, variations in data generation and preprocessing decisions, and similar sources of perturbation [24–26]. For example, upward-sloping NCPs typically have many small communities of good quality, so losing some communities via noise or some other perturbations has little effect on a realistic NCP. Naturally, whether a particular set of nodes achieves a local minimum is not robust to such modifications. In addition, one can easily construct pathological networks whose NCPs are not robust.

It is also important to consider the robustness of a network NCPs with respect to the use of conductance versus other measures of community quality. (Recall that many other measures have been proposed to capture the criteria that a good community should be densely-connected internally but sparsely connected to the rest of a network [5, 25].) Indeed, it has been shown that measures that capture both criteria of community quality (internal density and external sparsity) behave in a roughly similar manner to conductance-based NCPs, whereas measures that capture only one of the two criteria exhibit qualitatively different behavior, typically for rather trivial reasons [26].

Although the basic NCP that we have been discussing yields numerous insights about both small-scale and large-scale network structure, it also has important limitations. For example, an NCP gives no information on the number or density of communities with different community quality scores. (This contributes to the robustness properties of NCP with respect to perturbations of a network.) Accordingly, the communities that are revealed by an NCP need not be representative of the majority of communities in a network. However, the extremal features that are revealed by an NCP have important system-level implications for the behavior of dynamical processes on a network: they are responsible for the most severe bottlenecks for associated dynamical processes on networks [74].

Another property that is not revealed by an NCP is the internal structure of communities. Recall from Eq. (3) that the conductance of a community measures how well (relative to its size) that it is separated from the remainder of a network, but it does not consider the internal structure of a community (except for size and edge density). In an extreme case, a community with good conductance might even consist of several disjoint pieces. Recent work has addressed how spectral-based approximations to optimizing conductance also approximately optimize measures of internal connectivity [75].

We augment the information from basic NCPs with some additional computations. To obtain an indication of a community’s internal structure, we compute the internal conductance of the communities that form an NCP. The *internal conductance* $\phi_{\text{in}}(C)$ of a community C is

$$\phi_{\text{in}}(C) = \phi(G|_C), \quad (6)$$

where $G|_C$ is the subgraph of G induced by nodes in the community C . The internal conductance is equal to the conductance of the best partition into two communities of the network $G|_C$ viewed as a graph in isolation. Because a good community should be well-separated from the remainder of a network and also relatively well-connected internally, we expect good communities to have low conductance but high internal conductance. We thus compute the *conductance ratio*

$$\Phi(C) = \frac{\phi(C)}{\phi_{\text{in}}(C)} \quad (7)$$

to quantify this intuition. A good community should have a small conductance ratio, and thus we also plot so-called *conductance ratio profiles (CRPs)* [25] to illustrate how conductance ratio depends on community size in networks.

C. Our Application and Extension of NCPs

In this paper, we examine the small-scale, medium-scale, and large-scale community structure using conductance-based NCPs and CRPs. We employ three different methods, which we introduce in detail in Appendix B, for sampling an NCP: one based on local diffusion dynamics (the ACLCUT method), one based on a local spectral optimization (the MOV CUT method), and one based on geodesic distance from a seed node (the EGONET method). In each case, we find communities of different sizes, and we then plot the conductance of the best community for each size as a function of size.

An NCP provides a signature of community structure in a network, and we can thereby compare community structure across different networks. This helps one to discern which properties are attributable predominantly to network structure and which are attributable predominantly to choice of algorithms for community detection. Our approach of comparing community structures in networks using NCPs and CRPs is very general: one can of course follow a similar procedure with other community-quality diagnostics on the vertical axis, other procedures for community generation, and so on.

IV. EMPIRICAL RESULTS ON REAL NETWORKS

In this section, we present the results of our empirical evaluation of the small-scale, medium-scale, and large-scale community structure in our example networks.

A. Example Network Data Sets

We will examine six empirical networks in depth. They fall into three classes: coauthorship networks, Facebook networks, and voting similarity networks. For each class, we consider two networks of two different sizes.

- **Collaboration graphs.** The two (unweighted) coauthorship networks were constructed from papers submitted to the arXiv preprint server in the areas of general relativity and quantum cosmology (CA-GRQC) and Astrophysics (CA-ASTROPH). In each case, two authors are connected by an edge if they coauthored at least one paper, so a paper with k authors appears as a k -clique (i.e., a complete k -node subgraph) in the network. These network data are available as part of the Stanford Net-

work Analysis Package (SNAP), and they were examined previously in Refs. [24–26].

- **Facebook graphs.** The two (unweighted) Facebook networks are anonymized data sets that consist of a snapshot of “friendship” ties on one particular day in September 2005 for two United States (U.S.) universities: Harvard (FB-HARVARD1) and Johns Hopkins (FB-JOHNS55). They form a subset of the FACEBOOK100 data set from Refs. [11, 12]. In addition to the friendship ties, note that we possess node labels for gender and class year as well as numerical identifiers for student or some other (e.g., faculty) status, major, and high school.
- **Congressional voting graphs.** The two (weighted) Congressional voting networks represent similarities in voting patterns among members of the U.S. House of Representatives (US-HOUSE) and U.S. Senate (US-SENATE). Our construction follows prior work [9, 76]. In particular, we represent these two data sets as “multilayer” temporal networks [9, 77]. Each layer corresponds to a single two-year Congress, and edge weights within a layer represent the voting similarity between two legislators during the corresponding Congress. In layer s , this yields adjacency elements of $A_{ij}^{(s)} = \frac{1}{b_{ij}(s)} \sum_k \gamma_{ijk}$, where $\gamma_{ijk} = 1$ if both legislators voted the same way on the k^{th} bill, $\gamma_{ijk} = 0$ if they voted in different ways on that bill, $b_{ij}(s)$ is the number of bills on which both legislators voted during that Congress, and the sum is over bills. A tie between the same legislator in consecutive Congresses is represented by an inter-layer edge with weight ω [9]. (We use $\omega = 1$; the effect of changing ω has been investigated previously [76, 78].) We represent each multilayer voting network using a single “supra-adjacency matrix” (see Refs. [77, 79–81]) in which the different Congresses correspond to diagonal blocks and inter-layer edges correspond to off-block-diagonal terms in the matrix. Note that throughout this paper we treat the Congressional voting graphs at the level of this supra-adjacency matrix, without any additional labeling or distinguished treatment of inter- and intra-layer edges (cf. [9]).

We chose these three sets of networks because (as we will see in later sections) they have *very* different properties with respect to their large-scale versus small-scale community structures. We thus emphasize that, with respect to the topic of this paper, these six networks are representative of several broad classes of previously-studied networks: CA-GRQC and CA-ASTROPH are representative of the SNAP networks that were examined previously in Refs. [24–26]; both FB-HARVARD1 and FB-JOHNS55 (aside from a few very small communities in FB-HARVARD1) are representative of the FACEBOOK100 networks that were examined previously in Refs. [11, 12];

and US-HOUSE and US-SENATE give examples of networks (that are larger than the Zachary Karate Club and caveman networks) on which conventional notions of and algorithms for community detection have been validated successfully [9, 76].

In Table I, we provide summary statistics for each of the six networks. We give the numbers of nodes and edges in the largest connected component, the mean degree/strength ($\langle k_i \rangle$), the second-smallest eigenvalue (λ_2) of the normalized Laplacian matrix, and mean clustering coefficient ($\langle C_i \rangle$). We use the local clustering coefficient $C_i = \frac{1}{k_i(k_i-1)} \sum_{j,k} (\hat{w}_{ij}\hat{w}_{ik}\hat{w}_{jk})^{\frac{1}{3}}$, where $\hat{w}_{ij} = \frac{w_{ij}}{\max_{ij} w_{ij}}$, which reduces to the usual expression for local clustering coefficients in unweighted networks [82–84]. The high values for mean clustering coefficient in both the U.S. Congress and coauthorship networks are unsurprising, given how those networks have been constructed. However, the latter is noteworthy, as the coauthorship networks are much sparser than the Facebook networks.

Recall that the second smallest eigenvalue λ_2 of the normalized Laplacian provides a qualitative notion of connectivity that can be used to bound the mixing time of diffusion-based dynamics on networks [47] (where larger values of λ_2 imply that there are fewer bottlenecks to mixing) and that can also be used to partition a graph into communities [3, 65, 70] (where smaller values of λ_2 correspond to better communities). We show the values of λ_2 for our six networks in Table I. For comparison, we show in Fig. 4 a scatter plot of λ_2 versus the size of the network (i.e., the number of nodes in the network) for these six networks; for the remaining networks from the Stanford Network Analysis Project (SNAP) [85] (black circles) that were also studied in [24, 25]; and for the remaining 98 networks from the FACEBOOK100 data set (red stars) studied in [11, 12].

The first point to note about Fig. 4 is that λ_2 for nearly all of the FACEBOOK100 graphs is much larger than those for the two collaboration graphs and the two voting graphs. Figure 4 and previous empirical results (from Refs. [24, 25]) clearly demonstrate that the λ_2 values for the two collaboration graphs are representative of (and, in many cases, higher than) those of the other SNAP graphs studied empirically in Refs. [24, 25]. That is, nearly all of the networks have λ_2 values that are much smaller than those in the FACEBOOK100 graphs. This implies, in particular, that those graphs contain more substantial bottlenecks to mixing. (Note, though, that the value of λ_2 says nothing about the size or cardinality of the set of nodes that achieves the minimum.) In order to understand these differences, we study two networks from the FACEBOOK100 data set in detail: one (FB-JOHNS55) with a typical value of λ_2 and another (FB-HARVARD1) that is an “outlier,” in that it has the lowest value of λ_2 in the entire FACEBOOK100 data set. (The FB-Caltech36 network is the smallest network in the FACEBOOK100 data set—it has 762 nodes in its largest connected component (LCC)—and it has the largest value of λ_2 .)




FIG. 4. Scatter plot of the second smallest eigenvalue (λ_2) of the normalized Laplacian versus size of the network for: the networks from the SNAP data [85] that were studied in [24, 25]; all 100 networks in the FACEBOOK100 data set [11, 12]; and the two US-CONGRESS temporal networks [9, 76, 78].

The second point to note about Fig. 4 and Table I is that they suggest that FB-JOHNS55 (and possibly also FB-HARVARD1) are better connected than the other four networks, and that the connectivity properties of the two collaboration graphs and the two voting graphs (and perhaps also FB-HARVARD1) might be very similar. As we will see below, however, the situation is considerably more subtle.

In Fig. 5, we visualize the adjacency matrices of each of these networks using a sparsity-pattern (Spy) plot. We draw the nonzero entries of the adjacency matrix as black dots. The grayscale visualization in Fig. 5 is a result of coarsening the dpi-resolution and illustrates the density of connections in an area of the adjacency matrix. This yields a visualization comparable to the idealized block models in Fig. 1. The node order in a Spy plot is arbitrary and, by permuting the nodes, can sometimes yield visualizations that are suggestive of structural features in a network. For the coauthorship and Facebook networks, we use results from a single run of an implementation [86] of a Louvain-like heuristic [30] for modularity optimization to partition these networks into communities. We then sorted nodes by community assignment: we chose the order of communities manually to suggest potential large-scale structures. For the voting similarity networks, time provides a natural order for the nodes. We started with nodes from the 1st Congress and ended with the nodes from the 110th Congress. The small blocks on the diagonal are the individual Congresses, which are almost fully connected internally, and the off-diagonals result from the interlayer coupling between the same individuals from different Congresses.

While certainly not definitive, Fig. 5 suggests several hypotheses about the relationship between small-scale

| | Nodes | Edges | $\langle k \rangle$ | λ_2 | $\langle C \rangle$ | Refs. | Description |
|-------------|--------|-----------|---------------------|-------------|---------------------|-------------|---|
| CA-GRQC | 4 158 | 13 422 | 6.5 | 0.0019 | 0.56 | [24–26] | Coauthorship network: arXiv general relativity |
| CA-ASTROPH | 17 903 | 196 972 | 22.0 | 0.0063 | 0.63 | [24–26] | Coauthorship network: arXiv astrophysics |
| FB-JOHNS55 | 5 157 | 186 572 | 72.4 | 0.1258 | 0.27 | [11, 12] | Johns Hopkins Facebook network |
| FB-HARVARD1 | 15 086 | 824 595 | 109.3 | 0.0094 | 0.21 | [11, 12] | Harvard Facebook network |
| US-SENATE | 8 974 | 422 335 | 60.3 | 0.0013 | 0.50 | [9, 76, 78] | Network of voting patterns in U.S. Senate |
| US-HOUSE | 36 646 | 6 930 858 | 240.5 | 0.0002 | 0.58 | [9, 76, 78] | Network of voting patterns in U.S. House |

TABLE I. Six medium-sized networks. For each network, we show the number of nodes and edges in the largest connected component (LCC), the mean degree/strength ($\langle k_i \rangle$), the second-smallest eigenvalue (λ_2) of the normalized Laplacian matrix, the mean clustering coefficient ($\langle C_i \rangle$), prior references that used these networks, and a brief description.




FIG. 5. Sparsity-pattern (Spy) plots for the largest connected component of each of our six example networks. The coauthorship networks (CA-GRQC and CA-ASTROPH) and Facebook networks (FB-JOHNS55 and FB-HARVARD1) are arranged by communities that we obtained using an implementation [86] of a Louvain-like heuristic for modularity optimization [30]. For US-CONGRESS, we preserve the temporal order of the nodes starting with the first Congress in the top left and ending with the 110th Congress in the bottom right.

structure and the large-scale structure—and, in particular, between small communities and large communities—in these six networks. First, from Figs. 5c and 5f, it appears that the large-scale structure in US-SENATE and US-HOUSE corresponds to that of a “banded” matrix [87]. This banded structure is a result of the interlayer edges in these networks. Second, from Figs. 5a and 5d, it appears that CA-GRQC and CA-ASTROPH both have many small-scale communities. It appears that they have a large-scale structure that is roughly banded; but there also appear to be many “long-range” off-diagonal interactions between distant nodes in the depicted ordering. Third, from Figs. 5b and 5e, we observe that both FB-JOHNS55 and FB-HARVARD1 appear to have roughly 10 communities that are both relatively

large and relatively good.

From these visuals, it appears that nearly all of these communities have dense internal connections and sparse connections to other communities. Given the usual notion that communities are sets of nodes with denser connections among its constituent nodes than with the rest of the network, the visualizations in Fig. 5 appear to suggest that there might be interesting large-scale structure that might be exploitable in FB-JOHNS55 and FB-HARVARD1 but not in the other networks; and, in particular, that FB-JOHNS55 and FB-HARVARD1 seem to be examples of the case $\alpha_{11} \approx \alpha_{22} \gg \alpha_{12}$ illustrated in Fig. 1a.

The focus of the present investigation is to test the extent to which the above hypotheses about the relationship between small-scale structure and large-scale structure in these six networks is correct. As we have discussed, intuition like what we have illustrated in Fig. 5 is common in the development and validation of methods for community detection, so it is useful to delve into great depth on a set of networks to explore the connections between small-scale and large-scale connections in networks. As we will see in the next several sections, the situation is considerably more subtle than these figures (and commonly-employed intuition) might suggest. For example, with the exception of the small communities in CA-GRQC/CA-ASTROPH and the large-scale structure (i.e., the one-dimensional temporal ordering) in US-SENATE and US-HOUSE, these intuitive hypotheses about the relationship between the local structure and the global structure in these networks are not unambiguously supported by other evidence. Similarly, many communities that appear to be “good” based on the usual intuition and visualizations like that in Fig. 5 often are judged to be largely artifactual from the perspective of quantitative measures of community quality.

B. Network Community Profiles

We start by presenting our main results from using the ACLCUT method (see Figs. 6 and 7). One obtains similar insights about global structure using the MOV CUT (see

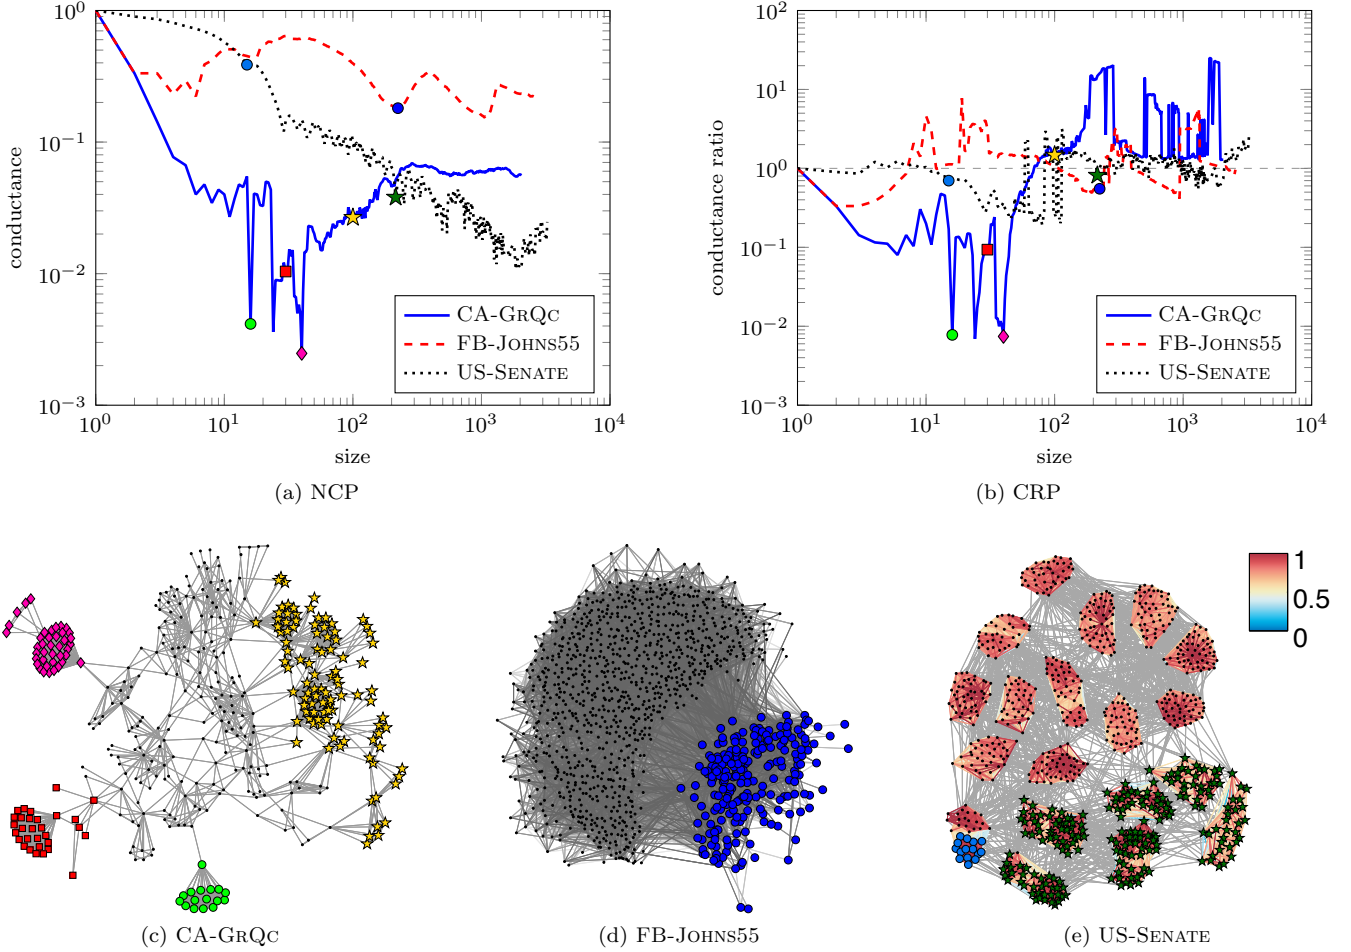


FIG. 6. NCP plots [in panel (a)] and conductance ratio profile (CRP) plots [in panel (b)] for CA-GRQC, FB-JOHNS55, and US-SENATE (i.e., the smaller network in each of the three pairs of networks from Table I) generated using the ACLCUT method. In panels (c)–(e), we show modified Kamada-Kawai [88] spring-embedding visualizations that emphasize community structure [89] of corresponding (color-coded) communities and their neighborhoods (2-neighborhood for CA-GRQC, a 1-neighborhood for FB-JOHNS55, and all Senates that have at least one Senator in common with the communities for US-SENATE). We find good small communities but no good large communities in CA-GRQC; some weak large-scale structure in FB-JOHNS55 that does *not* create substantial bottlenecks to the random-walk dynamics; and signatures of low-dimensional structure (i.e., good large communities but no good small communities) for US-SENATE, which results from the multilayer structure that encapsulates the network’s temporal properties.

Appendix C) and EGoNET (see Appendix D), although they can exhibit rather different local behavior.

In Fig. 6, we show the NCPs and CRPs for the smaller network from each of the three pairs of networks from Table I. In Fig. 7, we show the results for the corresponding larger networks. Note the logarithmic scale for both the vertical and horizontal axes in these figures as well as in subsequent NCP and CRP plots. Observe from Figs. 6a and 7a that the NCPs for networks of the same type are qualitatively similar, whereas NCPs for networks of different types have qualitatively distinct shapes.

- For the co-authorship networks **CA-GrQc** and **CA-AstroPh**, the NCPs have a mostly upward-sloping shape, except for the region with fewer than 100 nodes. We conclude that CA-GRQC and CA-

ASTROPH have good small (e.g., consisting of tens of nodes) communities, but they do not have good large (e.g., consisting of hundreds or thousands of nodes) communities. These results are consistent with the NCPs of LiveJournal from Fig. 2b and with the results of [24–26]. Additionally, the high values for the CRPs for the co-authorship networks (see Figs. 6b and 7b) for communities with hundreds or thousands of nodes reveals that these large communities are loosely connected collections of good, small communities. This feature is also visible in Fig. 6c, which shows selected communities and their neighborhoods for the CA-GRQC network.




FIG. 7. NCP plots [in panel (a)] and CRP plots [in panel (b)] for CA-ASTROPH, FB-HARVARD1, and US-HOUSE (i.e., the larger network in each of the three pairs of networks from Table I) generated using the ACLCUT method.

- For the Facebook networks **FB-Johns55** and **FB-Harvard1**, all of the communities at every size (except for two small “communities” with 5 and 10 nodes in FB-HARVARD1 [90]) have very large conductances (greater than 10^{-1}). This indicates that the communities in this network all have very poor community quality, in sharp contrast (though for different reasons) with both the co-authorship and voting networks. The essentially flat shape for the NCPs of the Facebook networks illustrate that these networks have strong expander-like properties (see Appendix A) and relatedly that there are no substantial bottlenecks to the rapid mixing of random walks on these networks. Both Facebook networks have noticeable dips in their NCPs at larger community sizes (about 220 and 1100 nodes for FB-JOHNS55, and about 1500 nodes for FB-HARVARD1), and the sets of nodes associated with each of these dips correlate strongly with self-reported demographic information [91].
- For the voting networks **US-Senate** and **US-HOUSE**, the NCP has a predominantly downward-sloping shape. This is characteristic of “low-dimensional” networks, in the sense that we described informally in Section II B. Informally, the reason for the downward-sloping shape is that US-SENATE and US-HOUSE consist of a low-dimensional structure that is evolving along a one-dimensional scaffolding (i.e., time), upon which the detailed structure of individual Congresses (i.e., a good partition that is nearly along party lines) is superimposed. (One can examine such structures by using smaller values of the interlayer coupling parameter; see Ref. [76].) This is consistent with previous results [92].

These results, which illustrate that community quality changes very differently with size in each of the three pairs of networks, also indicate that these three types of networks have very different properties with respect to large-scale versus small-scale community structure. Moreover, the qualitative similarity in behavior between the two networks in each pair suggests that the coarse behavior of an NCP (downward-sloping, upward-sloping, or flat) is indicative of large classes of networks and not an artifact of our particular choice of example networks. One obtains similar insights about global structure using the MOV CUT (see Appendix C) and EGONET (see Appendix D) methods, although they can exhibit rather different local behavior. We investigate these differences in local behavior in Section IV C.

C. Comparison of Results from ACLCUT, MOV CUT, and EGONET

The NCPs generated using either ACLCUT or MOV CUT (see Appendix C), and to a slightly lesser extent using EGONET (see Appendix D), have similar global features—i.e., they exhibit the same general trends and have dips at small size scales that correspond to nearly identical communities—indicating that we obtain a broadly similar picture of the large-scale community structure by using any of the methods. However, the detailed local behavior of the three methods can differ considerably. Such behavior depends sensitively on the choice of seed node, the choice of the parameters in the different methods, and the specific details of each method. In this section, we discuss the similarities and differences in the results from these methods. In this section, we only do calculations for the smaller networks from each of the three network pairs in Table I (but we

have observed similar results on the larger networks).

To compare different methods, we note that any meaningful difference between them should manifest itself as a difference in the rank order of nodes, as this determines the assignment of nodes to local communities. We quantify rank differences by computing the Spearman rank correlation [93] between the (exact for MOV CUT and approximate for ACL CUT) PPR and EgoRank ranking vectors. To make results from ACL CUT and MOV CUT comparable, we exploit the relation between γ and α (see Appendix B) to parametrize the MOV CUT method in terms of α . We also restrict all comparisons to the support of the corresponding approximate PPR vector that we obtained using the ACL CUT method. This induces an indirect dependency of the results from MOV CUT and EGONET on α and ϵ (in addition to the direct dependency of MOV CUT on α).

In Tables II–IV, we show the results of our calculations of Spearman rank correlations. For each of the three networks, we select 50 seed nodes by sampling uniformly without replacement. We then compute PPR vectors for these seed nodes using the ACL CUT and MOV CUT method for different values of the truncation parameter ϵ and teleportation parameter α , and we also compute the EgoRank vector for each of the seed nodes. Recall that *smaller* values of α correspond to *more local* versions of the procedures, but that *larger* values of ϵ correspond to *more local* versions of the procedures.

The ACL CUT and MOV CUT methods give very similar results for most of the 50 seed nodes in our sample, although (as discussed below) some seed nodes do yield noticeable differences. The two methods give the most similar results for FB-JOHNS55 (mean: 0.92, minimum: 0.43), whereas we find larger deviations in both CA-GRQC (mean: 0.85, minimum: -0.13) and US-SENATE (mean: 0.86, minimum: -0.44). Note that we calculated the mean, maximum, and minimum over all sampled seed nodes and parameter values.

Interestingly, the larger deviations between the two methods for CA-GRQC and US-SENATE occur at different values of the truncation parameter ϵ . For CA-GRQC (and, to a lesser extent, for FB-JOHNS55), we obtain the largest deviations for smaller values (e.g., $\epsilon = 10^{-6}$). For US-SENATE, however, we obtain the largest deviations for $\epsilon = 10^{-4}$. See the bold values in Tables II–IV. This is consistent with the very different isoperimetric properties of these three networks, as revealed by their NCPs, as well as with well-known connections between conductance and random walks.

There are two potential causes for the differences between the ACL CUT and MOV CUT method. First, there is a truncation effect, governed by the parameter ϵ , in approximating the PPR vector using the ACL CUT method. As ϵ becomes smaller, the approximation in ACL CUT becomes more accurate and this effect diminishes. Second, the two methods differ in the precise way that they use a seed vector to represent a seed node. Recall that the ACL CUT method uses an indicator vector \vec{s} to represent

a seed node i ; thus, we use $\vec{s}_i = 1$ whenever i is a seed node, and we set all other entries in that vector to 0. In contrast, the MOV CUT method projects the indicator vector onto the orthogonal complement of the strength vector to ensure that $\vec{s}^T D \vec{1} = 0$ (see Appendix B). This effect decreases as $\alpha \rightarrow 1$.

The larger deviations between the two methods occur for smaller values of ϵ in CA-GRQC and FB-JOHNS55; for these, the truncation effect is small, suggesting that the different way of representing a seed node is partially responsible for the difference between the results of the two methods for these networks. For larger values of ϵ (in particular, $\epsilon \geq 10^{-4}$), where the support of the approximate PPR vector from the ACL CUT method is small, the behavior of the two methods is very similar. Consequently, the differences in the choice of seed vector become more important for nodes that are “far away” from the seed node, in the sense that they are rarely visited by the personalized PageRank dynamics that underlie these methods. As a result, the “local NCPs” for the two methods in Figs. 8a and 9a are largely identical for small community sizes but diverge for large community sizes. (We use the term *local NCP* to refer to an NCP that we computed using only a single seed node without optimizing over the results from multiple seed choices; see Ref. [37] for details on the construction of local NCPs.)

For US-SENATE, the two methods behave almost identically for small ϵ (see Table IV), so we conclude that the different ways of representing a seed node have only a small effect on this network. However, the truncation effect is more pronounced in this network compared with CA-GRQC or FB-JOHNS55. This feature manifests as larger deviations between ACL CUT and MOV CUT in Table IV for large ϵ and small α (i.e., where the truncation has the strongest impact). The discrepancy occurs because the ACL CUT method initially pushes a large amount of probability to the interlayer neighbors of the seed node (i.e., to the same Senator in different Congresses). This probability does not diffuse to other nodes for sufficiently large values of ϵ .

In Figs. 8–10, we illustrate the results from Tables II–IV. In these figures, we plot the local NCPs for CA-GRQC, FB-JOHNS55, and US-SENATE for the seed nodes (from the sample of 50) that yield the highest and lowest mean Spearman rank correlation between the ACL CUT and MOV CUT methods. In these figures, we also include visualizations of example communities that we obtained from the ACL CUT and MOV CUT methods using a Kamada-Kawai-like spring-embedding visualization [89] of the k -ego-nets of these seed nodes.

From the visualizations of the local communities, it seems for CA-GRQC (see Fig. 8) and FB-JOHNS55 (see Fig. 9) that nodes included in local communities obtained from ACL CUT tend to be closer in geodesic distance than those obtained from MOV CUT to the seed node. (To see this, observe that red nodes tend to be larger than light blue nodes in the visualization of the k -neighborhoods.)

| | | α | | | | | | | | | | | | | | | |
|------------|-----------|----------|-------------|------|-------|-------------|-------|-------|--------------|------|-------|--------------|-------|-------|-------------|------|------|
| | | 0.6 | | | 0.7 | | | 0.8 | | | 0.9 | | | 0.99 | | | |
| | | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | |
| ϵ | 10^{-3} | max | 1.00 | 0.99 | 0.99 | 1.00 | 0.98 | 0.98 | 1.00 | 0.97 | 0.97 | 1.00 | 0.95 | 0.95 | 0.98 | 0.89 | 0.86 |
| | | mean | 0.98 | 0.78 | 0.77 | 0.98 | 0.76 | 0.73 | 0.98 | 0.72 | 0.68 | 0.97 | 0.68 | 0.62 | 0.91 | 0.60 | 0.48 |
| | | min | 0.92 | 0.26 | 0.23 | 0.91 | 0.18 | 0.14 | 0.94 | 0.21 | 0.15 | 0.85 | -0.01 | -0.05 | 0.74 | 0.25 | 0.06 |
| | 10^{-4} | max | 1.00 | 0.97 | 0.97 | 1.00 | 0.97 | 0.97 | 1.00 | 0.94 | 0.94 | 0.99 | 0.89 | 0.85 | 0.92 | 0.69 | 0.53 |
| | | mean | 0.99 | 0.74 | 0.72 | 0.98 | 0.73 | 0.68 | 0.97 | 0.70 | 0.64 | 0.95 | 0.63 | 0.54 | 0.89 | 0.51 | 0.36 |
| | | min | 0.96 | 0.16 | 0.10 | 0.91 | -0.05 | -0.19 | 0.84 | 0.35 | -0.02 | 0.88 | 0.43 | 0.25 | 0.85 | 0.32 | 0.18 |
| | 10^{-5} | max | 1.00 | 0.96 | 0.95 | 0.97 | 0.92 | 0.87 | 0.94 | 0.81 | 0.67 | 0.89 | 0.73 | 0.55 | 0.93 | 0.75 | 0.60 |
| | | mean | 0.91 | 0.74 | 0.58 | 0.89 | 0.69 | 0.51 | 0.85 | 0.65 | 0.42 | 0.78 | 0.62 | 0.33 | 0.84 | 0.63 | 0.36 |
| | | min | 0.24 | 0.21 | -0.20 | 0.42 | 0.30 | -0.10 | 0.42 | 0.39 | -0.13 | 0.25 | 0.43 | -0.10 | 0.49 | 0.44 | 0.05 |
| | 10^{-6} | max | 0.84 | 0.85 | 0.68 | 0.79 | 0.81 | 0.49 | 0.70 | 0.79 | 0.39 | 0.80 | 0.81 | 0.47 | 0.93 | 0.75 | 0.60 |
| | | mean | 0.62 | 0.72 | 0.25 | 0.57 | 0.69 | 0.17 | 0.51 | 0.69 | 0.12 | 0.51 | 0.72 | 0.13 | 0.85 | 0.63 | 0.37 |
| | | min | 0.01 | 0.57 | -0.36 | 0.07 | 0.52 | -0.30 | -0.06 | 0.50 | -0.27 | -0.13 | 0.57 | -0.32 | 0.52 | 0.46 | 0.08 |

TABLE II. Pairwise comparison of the three methods using the Spearman rank correlation between the rank vectors from the ACLCUT (A), MOV CUT (M), and EGO NET (E) methods for CA-GRQC. We use a uniform random sample of 50 nodes for each of several values for the teleportation parameter α and truncation parameter ϵ . We take the maximum, mean, and minimum over the seed nodes. Bold values highlight the largest deviations between ACLCUT and MOV CUT methods for a given value of α .

| | | α | | | | | | | | | | | | | | | |
|------------|-----------|----------|-------------|------|-------|-------------|------|-------|-------------|------|-------|-------------|------|------|-------------|------|------|
| | | 0.6 | | | 0.7 | | | 0.8 | | | 0.9 | | | 0.99 | | | |
| | | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | |
| ϵ | 10^{-3} | max | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | |
| | | mean | 1.00 | 0.85 | 0.84 | 1.00 | 0.82 | 0.81 | 0.99 | 0.79 | 0.79 | 0.98 | 0.78 | 0.77 | 0.98 | 0.78 | 0.75 |
| | | min | 0.95 | 0.61 | 0.56 | 0.94 | 0.54 | 0.47 | 0.80 | 0.55 | 0.53 | 0.80 | 0.55 | 0.49 | 0.77 | 0.40 | 0.39 |
| | 10^{-4} | max | 1.00 | 0.89 | 0.89 | 1.00 | 0.93 | 0.93 | 1.00 | 0.91 | 0.91 | 1.00 | 0.89 | 0.89 | 1.00 | 0.88 | 0.88 |
| | | mean | 0.98 | 0.48 | 0.47 | 0.97 | 0.45 | 0.44 | 0.96 | 0.43 | 0.41 | 0.95 | 0.41 | 0.37 | 0.94 | 0.38 | 0.33 |
| | | min | 0.81 | 0.06 | 0.05 | 0.78 | 0.05 | 0.04 | 0.74 | 0.07 | 0.06 | 0.69 | 0.07 | 0.07 | 0.69 | 0.07 | 0.00 |
| | 10^{-5} | max | 1.00 | 0.85 | 0.85 | 1.00 | 0.86 | 0.85 | 1.00 | 0.84 | 0.83 | 0.99 | 0.84 | 0.82 | 1.00 | 0.81 | 0.78 |
| | | mean | 0.98 | 0.58 | 0.54 | 0.97 | 0.59 | 0.54 | 0.97 | 0.63 | 0.57 | 0.96 | 0.61 | 0.54 | 0.95 | 0.53 | 0.46 |
| | | min | 0.90 | 0.24 | 0.18 | 0.91 | 0.26 | 0.17 | 0.91 | 0.23 | 0.14 | 0.89 | 0.22 | 0.07 | 0.88 | 0.18 | 0.04 |
| | 10^{-6} | max | 0.99 | 0.75 | 0.46 | 0.97 | 0.69 | 0.46 | 0.91 | 0.71 | 0.39 | 0.90 | 0.73 | 0.47 | 0.97 | 0.69 | 0.57 |
| | | mean | 0.79 | 0.41 | 0.20 | 0.75 | 0.45 | 0.13 | 0.72 | 0.56 | 0.20 | 0.77 | 0.62 | 0.32 | 0.88 | 0.59 | 0.41 |
| | | min | 0.57 | 0.23 | -0.07 | 0.49 | 0.24 | -0.06 | 0.43 | 0.32 | -0.02 | 0.49 | 0.32 | 0.01 | 0.60 | 0.26 | 0.07 |

TABLE III. Pairwise comparison of the three methods using the Spearman rank correlation between the ranking vectors from the ACLCUT (A), MOV CUT (M), and EGO NET (E) methods for FB-JOHNS55. We use a uniform random sample of 50 nodes for each of several values for the teleportation parameter α and truncation parameter ϵ . We take the maximum, mean, and minimum over the seed nodes. Bold values highlight the largest deviations between ACLCUT and MOV CUT methods for a given value of α .

If this observation holds more generally and is not just an artifact of the particular communities that we show in Figs. 8 and 9, then we should obtain higher Spearman rank correlations between ACLCUT and EGO NET than between MOV CUT and EGO NET. Indeed, Tables II–IV consistently show this effect for all choices of ϵ and α and for all three networks. Note that this effect is also present in US-SENATE, though it is less prominent in its k -neighborhood visualization than is the case for the

other two networks.

Figures 8–10 also reveal that the three networks look very different from a local perspective. For FB-JOHNS55 (see Fig. 9), both seed nodes that we considered result in reaching a large fraction of all nodes after just 2 steps. This is consistent with known properties of the full Facebook graph (circa 2012) of individuals connected by reciprocal “friendships.” For example, the mean geodesic distance between pairs of nodes of the Facebook graph

| | | α | | | | | | | | | | | | | | |
|------------|----------------|--------------|-------|-------|--------------|-------|-------|--------------|-------|-------|-------------|-------|-------|-------------|-------|-------|
| | | 0.6 | | | 0.7 | | | 0.8 | | | 0.9 | | | 0.99 | | |
| | | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E | A-M | A-E | M-E |
| ϵ | 10^{-3} max | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| | 10^{-3} mean | 0.65 | 0.66 | 0.40 | 0.68 | 0.71 | 0.30 | 0.74 | 0.77 | 0.46 | 0.74 | 0.85 | 0.57 | 0.80 | 0.80 | 0.56 |
| | 10^{-3} min | 0.02 | 0.50 | -0.05 | 0.38 | 0.32 | -0.24 | 0.46 | 0.53 | -0.07 | 0.50 | 0.68 | -0.02 | 0.51 | 0.62 | -0.00 |
| | 10^{-4} max | 1.00 | 0.93 | 0.87 | 1.00 | 0.92 | 0.77 | 1.00 | 0.93 | 0.82 | 1.00 | 0.89 | 0.75 | 0.93 | 0.81 | 0.61 |
| | 10^{-4} mean | 0.75 | 0.61 | 0.36 | 0.54 | 0.59 | 0.17 | 0.67 | 0.62 | 0.34 | 0.86 | 0.50 | 0.34 | 0.92 | 0.36 | 0.23 |
| | 10^{-4} min | -0.03 | -0.16 | -0.51 | -0.32 | -0.11 | -0.63 | -0.44 | 0.15 | -0.36 | 0.50 | -0.08 | -0.26 | 0.77 | -0.20 | -0.38 |
| 10^{-5} | max | 1.00 | 0.91 | 0.90 | 1.00 | 0.89 | 0.87 | 1.00 | 0.85 | 0.83 | 1.00 | 0.80 | 0.78 | 1.00 | 0.85 | 0.84 |
| | mean | 0.99 | 0.58 | 0.54 | 0.99 | 0.48 | 0.46 | 0.99 | 0.42 | 0.40 | 0.96 | 0.44 | 0.40 | 0.95 | 0.43 | 0.39 |
| | min | 0.86 | 0.03 | 0.01 | 0.88 | -0.02 | -0.02 | 0.88 | -0.09 | -0.13 | 0.80 | -0.17 | -0.20 | 0.80 | 0.07 | 0.01 |
| 10^{-6} | max | 1.00 | 0.84 | 0.84 | 1.00 | 0.84 | 0.84 | 1.00 | 0.84 | 0.85 | 1.00 | 0.87 | 0.86 | 1.00 | 0.96 | 0.96 |
| | mean | 0.94 | 0.54 | 0.48 | 0.98 | 0.55 | 0.50 | 0.99 | 0.62 | 0.60 | 1.00 | 0.69 | 0.68 | 0.99 | 0.91 | 0.89 |
| | min | 0.85 | 0.30 | 0.22 | 0.89 | 0.08 | 0.19 | 0.98 | 0.25 | 0.21 | 0.98 | 0.30 | 0.29 | 0.96 | 0.83 | 0.73 |

TABLE IV. Pairwise comparison of the three methods using the Spearman rank correlation between the ranking vectors from the ACLCUT (A), MOV CUT (M), and EGO NET (E) methods for US-SENATE. We use a uniform random sample of 50 nodes for each of several values for the teleportation parameter α and truncation parameter ϵ . We take the maximum, mean, and minimum over the seed nodes. Bold values highlight the largest deviations between ACLCUT and MOV CUT methods for a given value of α .

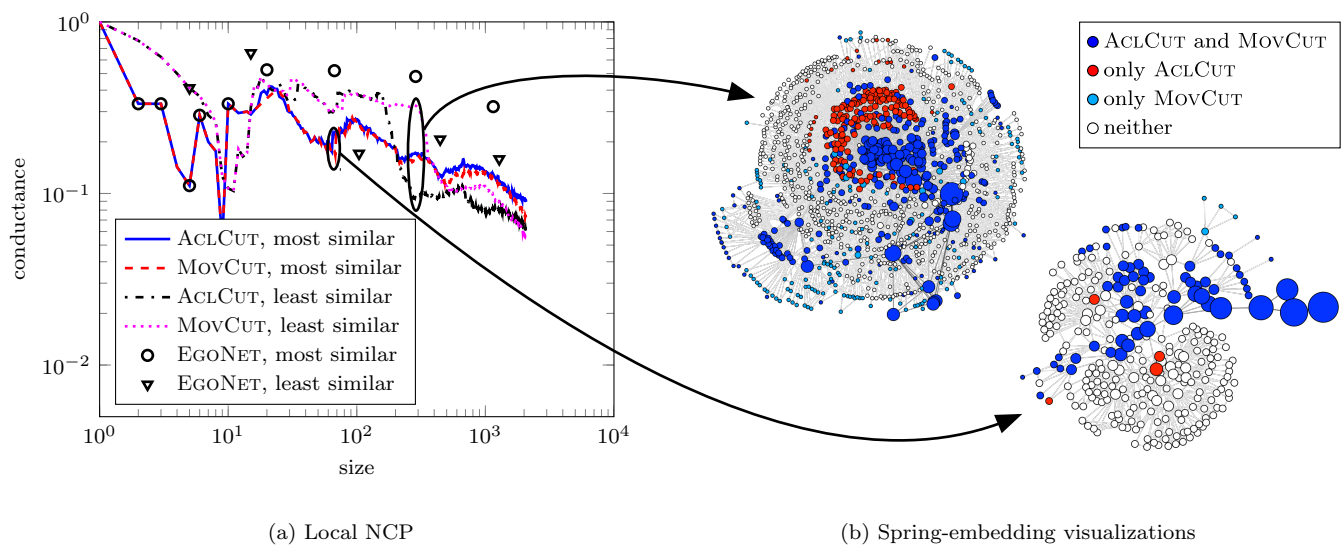


FIG. 8. CA-GRQC: (a) Local NCPs for the seed nodes (out of the 50 nodes that we sampled) with the highest and lowest mean Spearman correlation over the sampled parameter values. These NCPs highlight the difference in behavior for the two methods for large communities. (b) Kamada-Kawai-like spring-embedding visualization [89] of (bottom right) the 9-neighborhood of the seed node with the smallest difference between the two methods and (top left) the 6-neighborhood of the seed node with the largest difference. In these two visualizations, the node size decreases as a function of geodesic distance from the seed node. We color the nodes according to whether they belong to the local community that we obtained using the ACLCUT method, the one we obtained using the MOV CUT method, or both methods.

is very small: it was recently estimated by Facebook’s Data Team and their collaborators to be about 4.74 [94]. Additionally, as reported by Facebook’s Data Team, one can view Facebook as a collection of ego networks that have been patched together into a network whose global structure is sparse [44] (and such structure is an important motivation for the locally-biased notion of commu-

nity structure that we advocate in this paper).

For CA-GRQC, we obtain very different neighborhoods starting from our two different seed nodes. The node that exhibits the largest difference in behavior for both the ACLCUT and MOV CUT methods appears to be better connected in the network in the sense that the k -neighborhood (for any k until saturation occurs) is much




FIG. 9. FB-JOHNS55: (a) Local NCPs for the seed nodes (out of the 50 nodes that we sampled) with the highest and lowest mean Spearman correlation over the sampled parameter values. These NCPs highlight the difference in behavior for the two methods for large communities. (b) Kamada-Kawai-like spring-embedding visualization [89] of the 2-neighborhoods of both seed nodes. The one with the smallest difference in the bottom right and the one with the largest difference in the top left. In these two visualizations, the node size decreases as a function of geodesic distance from the seed node, but they appear in at least one of the local communities. We color the nodes according to whether they belong to the local community that we obtained using the ACLCUT method, the one we obtained using the MOV CUT method, or both methods.




FIG. 10. US-SENATE: (a) Local NCP for the seed nodes out of the 50 nodes sampled with the highest and lowest average Spearman correlation over sampled parameter values. These NCPs highlight the difference in behavior for the two methods for large communities. (b) Kamada-Kawai-like spring-embedding visualization [89] of the 3-neighborhoods of both seed nodes. The one with the smallest difference in the bottom right and the one with the largest difference in the top left. In these two visualizations, the node size decreases as a function of geodesic distance from the seed node. We color the nodes according to whether they belong to the local community that we obtained using the ACLCUT method, the one we obtained using the MOV CUT method, or both methods.

larger than that of the node that showed the smallest difference. (That is, it is more in the “core” than in the “periphery” of the nested core-periphery structure of Refs. [24, 25].) We observe a similar phenomenon for FB-

JOHNS55 and US-SENATE. Furthermore, its 1-ego-net and 2-ego-net are highly clustered, in the sense that they contain many closed triangles. For the seed node that showed the smallest difference between the ACLCUT and

MOVCUT methods, we need to consider the 6-ego-net (which has 20 nodes) to obtain a network of similar size to the 2-ego-net for the seed node with the largest difference (which has 15 nodes). In the case of the seed node in our sample that showed the least difference between the two methods, even the 6-ego-net appears rather tree-like; it contains few closed triangles and no larger cliques.

For US-SENATE, the 1-neighborhood of any seed node contains only the node itself and those corresponding to the same Senator in different Congresses [95]. As one begins to consider nodes that are further away, one first reaches corresponding Senators in other Congresses before reaching other Senators with similar voting patterns from the same Congress. This behavior of the EGONET method contrasts with the (PageRank-based) ACLCUT and MOVCUT methods, which tend to initially select all Senators from one Congress before reaching Senators from other Congresses.

D. Meso-Scale Structure

From the perspective of the locally-biased community-detection methods that we use in this paper, one can view intermediate-sized (i.e., meso-scale) structures in networks as arising from collections of local features—e.g., via overlaps of local communities that one obtains algorithmically using locally-biased dynamics such as those that we consider. Such local features depend not only on the network adjacency matrix but also on the dynamical process under study, the initial seed(s) from which one is viewing a network, and the locality parameters of the method (which corresponds to the dynamical process) that determine how locally one is viewing the network. Although a full discussion of the relationship between local structure and meso-scale structure and global structures is beyond the scope of this paper, here we provide an initial example of such results.

To try to visualize meso-scale and global network structures that we obtain from the local communities that we identify, we define an $n \times n$ association matrix \tilde{A} (where n is again the number of nodes in the network), which encodes pairwise relations between nodes based on a sample of local communities. For a given sample \mathcal{S} of local communities (obtained, e.g., by running a given method with many seed nodes and values of a locality parameter), the entries of the association matrix are given by the number of times that a pair of nodes appear together in a local community, normalized by the number of times either of them appeared. That is, the elements of the association matrix are

$$\tilde{A}_{ij} = \frac{|\{S \in \mathcal{S} : i \in S \text{ and } j \in S\}|}{|\{S \in \mathcal{S} : i \in S \text{ or } j \in S\}|}. \quad (8)$$

Our procedure for extracting global network structure from a sampled set of communities is similar in spirit to computing association (or “co-classification”) matrices that have been constructed from sampling a land-

scape of the modularity objective function [96], and one can in principle analyze these matrices further using the same methods. The additional normalization in our definition of association matrices is necessary to correct for the oversampling of large communities relative to small communities (which results from sampling nodes uniformly at random). At first glance, association matrices computed by sampling a modularity landscape appear to reveal much clearer community structure in these networks than what we obtain by sampling local communities. However, this is largely an artifact of the well-known resolution limit of modularity optimization [35]. One can mitigate this effect by using one of the multi-resolution generalizations of modularity [97, 98] to sample the modularity landscape across different values of the resolution parameter. This yields association matrices that are similar in appearance to the ones that we obtain by sampling local communities.

To visualize the association matrices in a way that reveals global network structure, it is important to find a good node order. We found the sorting method suggested in Ref. [96] to be impractically slow for the networks that we study. Instead, we sort the nodes based on the optimal leaf ordering [99] for the average-linkage hierarchical clustering tree of the association matrix. (For US-SENATE, we do this procedure within a given Congress, and we then use the natural temporal ordering to define the inter-Congressional ordering.)

In addition, to see small-scale structure using samples \mathcal{S} obtained from MOVCUT, we use a community-size parameter c that limits the volume of the resulting community based on the desired correlation with the seed vector. In this paper, we use $c \in \{10^i : i = 1, \dots, 5\}$. See Ref. [37] for details. We summarize our results in Figs. 11–14.

In Fig. 11, we show the result of applying this procedure with communities that we sampled using the ACLCUT, MOVCUT, and EGONET methods. In each case, we keep only the best conductance community for each sampled ranking vector. The most obvious feature of the visualizations in Fig. 11 is that—except for US-SENATE, for which there is a natural large-scale global structure defined by the one-dimensional temporal ordering—the visualizations are much more complicated than any of the idealized structures in Fig. 1 (which suggests that the visualizations might be revealing at least as much about the inner workings of the visualization algorithm as about the networks being visualized). The structures in Fig. 1 are trivially interpretable, whereas those in real networks (e.g., as illustrated in Fig. 11) are extremely messy and very difficult to interpret. In the paragraphs below, we will discuss the structural features in Fig. 11 in more detail.

For CA-GRQC (see Fig. 12 as well as Fig. 11), we observe many small communities that are composed of about 10–100 nodes. These communities, which correspond to the dark red blocks along the diagonal (see the inset in Fig. 11a), are responsible for the dips in the




FIG. 11. Visualizations of association matrices for CA-GRQC, FB-JOHN55, and US-SENATE illustrate how meso-scale and global structures emerge from the superposition and overlap of many local communities. See the main text for a description of how we construct the association matrices. For each of the three networks, we generate the subfigures using the same three sampling procedures that we use to generate the NCPs: we use ACLCUT for panels (a)–(c), we use MOV CUT for panels (d)–(f), and we use EGO NET for panels (g)–(i).

NCPs (see Figs. 6a, 16a, and 18a) for this network. However, these small communities do not combine to form large communities, which would result in large diagonal blocks in the association matrices. Instead, the small communities appear to amalgamate into a single large

block (or “core”). In Fig. 12, we aim to make this observation more intuitive by showing how the local communities for three different seed nodes spread through the network as we change the resolution, i.e., the locality bias parameter. We construct the weighted network




FIG. 12. Visualization of global structure in CA-GRQC. We constructed the network layout by weighting each edge using the corresponding entry of the association matrix for the ACLCUT method (Fig. 11a). We then applied the spring embedding visualization algorithm [89] to the resulting weighted network. For ease of visualization, we only plot edges with weight larger than the mean edge weight. Colored nodes correspond to local communities for three different seed nodes. Nodes that are a member of more than one community are drawn in a mixed color (e.g., blue and yellow become green; and blue, yellow, and red become blackish, in panel (d)). As we decrease the resolution parameter ϵ , the different communities first explore local structure before merging and each covering most of the network, in panel (d).

$\tilde{G} = (V, E, \tilde{w})$ shown in Fig. 12 from the unweighted CA-GRQC network $G = (V, E)$ using the association matrix for the ACLCUT method (Fig. 11a). We assign each edge $(i, j) \in E$ a weight based on the corresponding entry of the association matrix, i.e., $\tilde{w}_{ij} = \tilde{A}_{ij}$ if $(i, j) \in E$ and $\tilde{w}_{ij} = 0$ otherwise. Based on our earlier results with the slowly-increasing NCP, as well as previous results in Refs. [24–26], we interpret these features shown in Fig. 12 in terms of a nested core-periphery structure, in which the network periphery consists of relatively good communities and the core consists of relatively densely connected nodes.

For FB-JOHNS55 (see Fig. 13 as well as Fig. 11), we observe two relatively large communities, which correspond to the two large diagonal blocks in Figs. 11b and 11e and which underlie the dips in the NCPs in Figs. 6a and 16a. Note, however, from the scale of the vertical axis

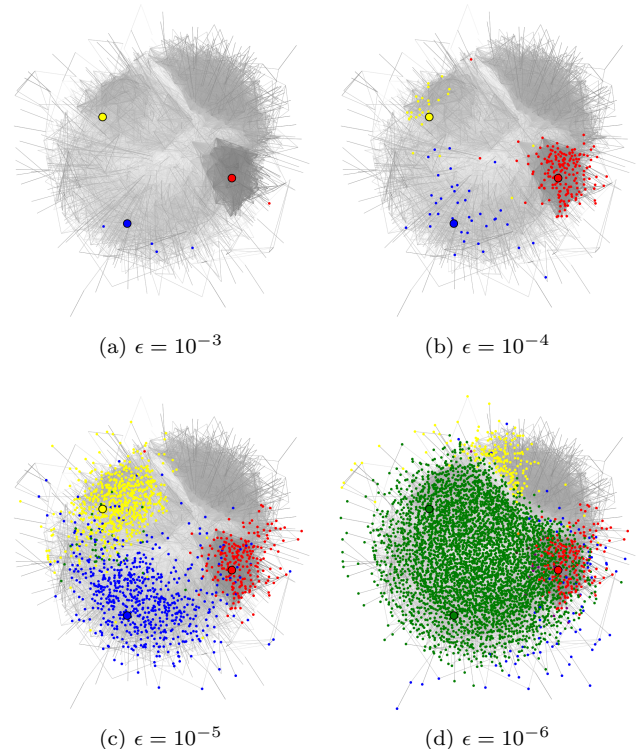


FIG. 13. Visualization of global structure in FB-JOHNS55. We constructed the network layout by weighting each edge using the corresponding entry of the association matrix for the ACLCUT method (Fig. 11b) and then applying the same procedure as in Fig. 12. Colored nodes correspond to local communities for three different seed nodes. Note the difference in behavior for the red community versus the blue and yellow communities. The blue and yellow communities gradually spread as we decrease ϵ , and they eventually merge to cover a large part of the network. However, the red community quickly spreads initially as we decrease ϵ but then remains localized as we decrease ϵ further.

in Figs. 6a and 16a that the community quality of these communities is very low, so one should actually construe the visualization in Figs. 11b and 11e as highlighting a low-quality community that is only marginally better than the other low-quality communities that are present in that network. Based on this visualization as well as our earlier results, the remainder of FB-JOHNS55 does not appear to have much community structure (at least based on using the conductance diagnostic to measure internal versus external connectivity). However, there do appear to be some remnants of highly overlapping communities that one could potentially identify using other methods (e.g., the one in Ref. [36]). The EGO NET method (see Fig. 11h) is unable to resolve not only these small communities but also the larger low-quality communities. Figure 13 shows how the local communities for two seed nodes that do not belong to one of the two large communities slowly spread and eventually merge (blue and yellow nodes), whereas the red community (which corre-




FIG. 14. Visualization of global structure in US-SENATE. We constructed the network layout by reweighting each edge using the corresponding entry of the association matrix for the ACLCUT method (Fig. 11c) and then applying the same procedure as in Figs. 12 and 13. Colored nodes correspond to local communities for three different seed nodes. The spreading behavior of the different local communities largely follows the temporal structure of the network.

sponds to the smaller of the two communities) is quickly identified and remains separate from the other communities.

For US-SENATE (see Fig. 14 as well as Fig. 11), we clearly observe the signature of temporal-based community structure at a large size scale. See Figs. 11c, 11f, and 11i. Using ACLCUT and MOV CUT, we also obtain partitions at the scale of individual Congresses (see the insets in Figs. 11c and 11f), which sometimes split into two or occasionally three individual communities. These latter partitions have been discussed previously in terms of polarization between parties [9, 92, 100]. Because we fixed the temporal order of Congresses for US-SENATE and only sort Senators within the same Congress, this visualization reveals communities within each Senate as well as more temporally-disparate communities. In particular, for the EGONET method, this ordering introduces a checkerboard pattern that correspond to temporal communities that contain Senators from several Congresses. Figure 14 clearly shows that this temporal structure also dominates the behavior of local communities for individual seed nodes.

An important point from these visualizations is that,

for both CA-GRQC and FB-JOHNS55, the meso-scale and large-scale structures that result from the superposition of local communities does *not* correspond particularly well to intuitive good-conductance communities. Relatedly, it also does *not* correspond particularly well to an intuitive low-dimensional structure or a nearly decomposable block-diagonal matrix of community assignments (see our illustration in Fig. 1a), one or both of which are often assumed (typically implicitly) by many global methods for algorithmically detecting communities in networks [5, 6, 37, 101, 102]. Of the networks that we investigate, only the temporal structure in US-SENATE (as well as in US-HOUSE, which is a related temporally-dominant network) closely resembles such an idealization. This is reflected clearly in its downward-sloping NCP (see Figs. 6a, 16a, and 18a) and in the visualizations in Fig. 11.

Instead, in the other (e.g., collaboration, Facebook, and many many other realistic [24, 25]) networks, community structure as a function of size is much more subtle and complicated. Fortunately, our locally-biased perspective provides one means to try to resolve such intricacy. By averaging over results from different seed nodes, a local approach like ours leads naturally to the presence of strongly overlapping communities. Overlapping community structure has now been studied for several years [103–105], and recent observations continue to shed new light on the ubiquity of community overlap [36]. Overlap of communities in networks is a pervasive phenomenon [36, 106]; and our expectation is that most large realistic networks have communities with significant overlap, rather than merely a small amount of overlap that would amount to a small perturbation of the idealized, nearly decomposable communities in Fig. 1a. Additionally, such overlaps imply that larger communities tend to have lower quality in terms of their internal versus external connectivity (i.e., in terms of how much they resemble the intuitive communities that many researchers know and love) than smaller communities—in agreement with our empirical results on both the collaboration networks and Facebook networks, but in strong disagreement with popular intuition. In these latter cases, recent work that fits related networks with upward-sloping NCPs to hierarchical Kronecker graphs resulted in parameters that are consistent with the core-periphery structure that we illustrated in Fig. 1b [107].

V. EMPIRICAL RESULTS ON SYNTHETIC BENCHMARKS

Synthetic benchmark networks with a known, planted community structure can be helpful for validating and gaining a better understanding of the behavior of community-detection algorithms. For such an approach to be optimally useful, it is desirable for the synthetic benchmarks to reproduce relevant features of real networks with community structure; and it is challenging




FIG. 15. NCPs of LFR synthetic benchmark networks [108] with $n = 10000$ nodes. Colors correspond to different values of the mixing parameter μ . Our choices for the mean degree $\langle k \rangle$, maximum degree k_{\max} , exponent of the degree distribution τ_1 , exponent of the community size distribution τ_2 , minimum community size c_{\min} , and maximum community size c_{\max} correspond to the ones used in Refs. [28, 109] to benchmark community-detection algorithms.

to develop good benchmarks that reproduce community structure and other structural properties of medium-sized and larger realistic networks. An extremely popular—and in some ways useful—family of benchmark networks that aims to reproduce some features of real networks are the so-called *LFR (Lancichinetti-Fortunato-Radicchi) networks* [108, 110]. By design, LFR networks have power-law degree distributions as well as power-law community-size distributions, they are unweighted, and they have non-overlapping planted communities. Motivated by our empirical results on networks constructed from real data, we also apply our methods to LFR networks to test the extent to which they are able to reproduce the three classes of NCP behavior (upward-sloping, flat, and downward-sloping) that we have observed with real networks.

To parametrize the family of LFR networks, we specify its power-law degree distribution using its exponent τ_1 , mean degree $\langle k \rangle$, and maximum degree k_{\max} . Similarly, we specify its power-law community size distribution using its exponent τ_2 , minimum community size c_{\min} , and maximum community size c_{\max} , with the additional constraint that the sum of community sizes should equal the size of the network n . Furthermore, we specify the strength of community memberships using a mixing parameter μ , where each node shares a fraction $1 - \mu$ of its edges with nodes in its own community. A simple calculation shows that this definition of the mixing parameter implies that each community in the planted partition has conductance μ (up to rounding effects).

To construct a network with these parameters, we sample n degrees from the degree distribution and sample community sizes from the community size distribution. We then assign nodes to communities uniformly at random, with the constraint that a node cannot be assigned to a community that is too small for the node to have the correct mixing-parameter value. We then construct inter-community and intra-community edges sep-

arately by connecting the corresponding stubs (i.e., ends of edges) uniformly at random. We use the implementation by Lancichinetti [111] to generate LFR networks.

In Fig. 15, we show representative NCPs for LFR networks for three choices of parameters for the degree distribution and community-size distribution that have been used previously to benchmark community-detection algorithms [28, 108, 109]. (We generated the results presented in Fig. 15 using the ACLCUT method, but we obtain nearly identical NCPs using the MOV CUT method.) The three subfigures demonstrate that all three parameter choices yield networks with similar NCPs. In particular, we observe that—above a certain critical size—the best communities have comparable quality, as a function of increasing size. Depending on the particular parameter values, this can be of similar quality to or somewhat better than that which would be obtained by, e.g., a vanilla (not extremely sparse) ER random graph, across all larger size scales. That is, above the critical size, the NCP is approximately flat. Increasing the topological mixing parameter μ in the LFR network generative mechanism at first shifts the entire NCP upwards because the number of inter-community edges increases. For $\mu \approx 1$, it levels off to the characteristic flat shape for an NCP of a network generated from the configuration model of random graphs.

Importantly, the behavior for the LFR benchmark networks from Ref. [108] that we illustrate in Fig. 15 does *not* resemble the NCPs for any of the real-world networks in either the present paper or in Ref. [24, 25]. In addition, we have been unable to find parameter values for which the qualitative properties of realistic NCPs—in particular, a relatively gradually upward-sloping NCP—are reproduced, which suggests that the community structure generated by the LFR benchmarks is *not* realistic in terms of its size-resolved properties.

To verify that this behavior is not an artifact of the particular choices of parameters shown in

Fig. 15, we sampled sets of parameters uniformly at random with $n \in \{1000, 10\,000, 50\,000\}$, $\tau_1, \tau_2 \in \{-1, -2, \dots, -5\}$, $\langle k \rangle \in \{10, 11, \dots, 100\}$, $k_{\max} \in \{\langle k \rangle, \langle k \rangle + 1, \dots, 250\}$, $c_{\min} \in \{10, 11, \dots, 250\}$, and $c_{\max} \in \{\max(c_{\min}, k_{\max}), \dots, 250\}$. The aggregate trends of the NCPs for the LFR benchmark networks with the different parameters we sample are similar to and consistent with the results shown in Fig. 15. Hence, although the LFR benchmark networks are useful as tests for community-detection techniques, our calculations suggest that they are unable to reproduce a fundamental feature of many real networks with respect to variation in community quality (and, in particular, worsening community quality) as a function of increasing community size.

Based on our empirical observations, our locally-biased perspective on community detection suggests a natural approach to determine whether synthetic benchmarks possess small-scale, medium-scale, and large-scale community structure that resembles that of large realistic networks: namely, a family of synthetic benchmark networks ought to include parameter values that generate networks with (robust) upward-sloping, flat, and downward-sloping NCPs (as observed in Figs. 2a and 6a).

VI. CONCLUSIONS AND DISCUSSION

In this paper, we have conducted a thorough investigation of community quality as a function of community size in a suite of realistic networks, and we have reached several conclusions with important implications for the investigation of realistic medium-sized to large-scale networks. Our results build on previous work on using network community profiles (NCPs) to study large-scale networks [24–26]. In this paper, we have employed a wider class of community-identification procedures, and we have discovered a wider class of community-like behaviors (as a function of community size) in realistic networks than what had been reported previously in the literature [112]. In addition, using NCPs, we have discovered that the popular LFR synthetic benchmark networks, which are often used to validate community-detection algorithms—and which are the most realistic synthetic benchmark networks that have been produced to test methods for community detection [113]—exhibit behavior that is markedly different from many realistic networks. Our result thus underscores the importance of developing realistic benchmark graphs whose NCPs are qualitatively similar to those of real networks. Taken together, our empirical results yield a much better understanding of realistic community structure in large realistic networks than was previously available, and they provide promising directions for future work. More generally, because our approach for comparing community structures in networks (using NCPs and conductance ratio profiles) is very general—e.g., one can follow an analogous procedure with other community-quality diagnos-

tics, other procedures for community generation, etc.—our locally-biased and size-resolved methodology is an effective way to investigate size-resolved meso-scale network structures much more generally.

The main conclusion of our work is that community structure in real networks is much more intricate than what is suggested by the block-diagonal assumption that is (either implicitly or explicitly) made by most community-detection methods (including ones that allow overlapping communities [103]) and when using the synthetic benchmark networks that have been developed to test those methods. Community structure interplays with other meso-scale features, such as core-periphery structure [36, 62, 64], and investigating only community structure without consideration of other structures can lead to misleading results. A local perspective on community detection, like the one that we have advocated in the present paper, allows pervasive community overlap in a natural way—which is an important feature to capture when considering real social networks. Additionally, the large-scale consensus community structure that we obtain subsequently by “pasting together” local communities is not constrained to resemble a global block-diagonal structure. This is a key consideration in the study of meso-scale structures in real networks.

Although most algorithmic methods for community detection take a different approach from ours, the observation that network community structure depends not only on the network structure per se but also on the dynamical processes that take place on a network and the initial conditions (i.e., seed node or nodes) for those processes, is rather traditional in many ways. Recall, for example, Granovetter’s observation that a node with many weak ties is ideally suited to initialize a successful social contagion process [114]. Our perspective also meshes better than global ones with real-life experience in our own networks. Both of these observations underscore our point that whether particular network structures form bottlenecks for a dynamical process depends not only on the process itself but also on the initial conditions of that process.

More generally, one might hope that our size-resolved and locally-biased perspective on community detection can be used to help develop new diagnostics that complement widely-used and intuitive concepts such as closeness centrality, betweenness centrality, and the many other existing global notions. These will be of particular interest for investigating large networks—or even modestly-sized networks such as those that we have considered—where traditional algorithmic and visualization methods have serious difficulties. Because the study of meso-scale structure in networks is important for understanding how local and small-scale properties of a network interact with global or large-scale properties, we expect that taking a locally-biased perspective on community detection and related problems will yield interesting and novel insights on these and related questions.

Appendix A: Expander Graphs

In this section, we provide a brief introduction to the concept of an *expander graph* (or, more simply, an *expander*) [115]. Essentially, expanders are graphs that are very well-connected and thus do not have any good communities (when measured with respect to diagnostics such as conductance). Because our empirical results indicate that many large social and information networks are expanders—at least when viewed at large size-scales—it is useful to review basic properties about expander graphs. Although most of the technical aspects of expander graphs are beyond the scope of this paper, Ref. [116] provides an excellent overview of this topic.

Let $G = (V, E)$ be a graph, which we assume for simplicity is undirected and unweighted. For the moment, we assume that all nodes have the same degree d (i.e., G is d -regular). For $S_1, S_2 \subset V$, the set of edges from S_1 to S_2 is then

$$E(S_1, S_2) = \{(u, v) : u \in S_1, v \in S_2, (u, v) \in E\}. \quad (\text{A1})$$

In this case, the number $|S|$ of nodes in S is a natural measure of the size of S . Additionally, the quantity $|E(S, \bar{S})|$, which indicates the number of edges that cross between S and \bar{S} , is a natural measure of the size of the boundary between S and \bar{S} .

We also define the *edge expansion of a set of nodes* $S \subset V$ as

$$h(S) = \frac{|E(S, \bar{S})|}{|S|}, \quad (\text{A2})$$

in which case the *edge expansion of a graph* G is the minimum edge expansion of any subset (of size no greater than $n/2$) of nodes:

$$h(G) = \min_{S \subset V: |S| \leq \frac{n}{2}} h(S). \quad (\text{A3})$$

A sequence of d -regular graphs $\{G_i\}_{i \in \mathbb{N}}$ is a *family of expander graphs* if there exists an $\epsilon > 0$ such that $h(G_i) \geq \epsilon$ for all $i \in \mathbb{N}$. Informally, a given graph G is an expander if its edge expansion is large.

As reviewed in Ref. [116], one can view expanders from several complementary viewpoints. From a combinatorial perspective, expanders are graphs that are highly connected in the sense that one has to sever many edges to disconnect a large part of an expander graph. From a geometric perspective, this disconnection difficulty implies that every set of nodes has a relatively very large boundary. From a probabilistic perspective, expanders are graphs for which the natural random-walk process converges to its limiting distribution as rapidly as possible. Finally, from an algebraic perspective, expanders are graphs in which the first nontrivial eigenvalue of the Laplacian operator is bounded away from 0. (Because we are talking here about d -regular graphs, note that this statement holds for both the combinatorial Laplacian and the normalized Laplacian.) In addition, constant-degree

(i.e., d -regular, for some fixed value of d) expanders are the metric spaces that (in a very precise and strong sense [116]) embed least well in low-dimensional spaces (such as those discussed informally in Section II B). All of these interpretations imply that smaller values of expansion correspond more closely to the intuitive notion of better communities (whereas larger values of expansion correspond, by definition, to better expanders.)

Note the similarities between Eq. (A2) and Eq. (A3), which define expansion, with Eq. (3) and Eq. (4), which define conductance. These equations make it clear that the difference between expansion and conductance simply amounts to a different notion of the size (or volume) of sets of nodes and the size of the boundary (or surface area) between a set of nodes and its complement. This difference is inconsequential for d -regular graphs. However, because of the deep connections between expansion and rapidly-mixing random walks, the latter notion (i.e., conductance) is much more natural for graphs with substantial degree heterogeneity. The interpretation of failing to embed well in low-dimensional spaces (like lines or planes) is not as extremal in the case of conductance and degree-heterogeneous graphs as it is in the case of expansion and degree-homogeneous graphs; but the interpretations of being well-connected, failing to provide bottlenecks to random walks, etc. all hold for conductance and degree-heterogeneous graphs such as those that we consider in the main text of the present paper. Accordingly, it is insightful to interpret our empirical results on small-scale versus large-scale structures in networks should be in light of known facts about expanders and expander-like graphs.

Appendix B: Community Quality, Dynamics on Graphs, and Bottlenecks to Dynamics

In this section, we describe in more detail how we algorithmically identify possible communities in graphs. Because we are interested in local properties and how they relate to meso-scale and global properties, we take an operational approach and view communities as the output of various dynamical processes (e.g., diffusions or geodesic hops), and we discuss the relationship between the output of those procedures to well-defined optimization problems. The idea of using dynamics on a network has been exploited successfully by many methods for finding “traditional” communities (of densely connected nodes) [9, 32, 53, 117–120] as well as for finding sets of nodes that are related to each other in other ways [48, 54, 117, 121, 122].

In this paper, we build on the idea that random walks and related diffusion-based dynamics, as well as other types of local dynamics (e.g., ones, like geodesic hops, that depend on ideas based on egocentric networks), should get “trapped” in good communities. In particular, we consider the following three dynamical methods for community identification.

1. Dynamics Type 1: Local Diffusions (the “ACL CUT” method).

In this procedure, we consider a random walk that starts at a given seed node s and runs for some small number of steps. We take advantage of the idea that if a random walk starts inside a good community and takes only a small number of steps, then it should become trapped inside that community. To do this, we use the locally-biased personalized PageRank (PPR) procedure of Refs. [123, 124]. Recall that a PPR vector is implicitly defined as the solution of the equation

$$\text{pr}(\alpha, \vec{s}) = \alpha D^{-1} A \text{pr}(\alpha, \vec{s}) + (1 - \alpha) \vec{s}, \quad (\text{B1})$$

where $1 - \alpha$ is a “teleportation” probability and \vec{s} is a seed vector. From the perspective of random walks, evolution occurs either by the walker moving to a neighbor of the current node or by the walker “teleporting” to a random node (e.g., determined uniformly at random as in the usual PageRank procedure, or to a random node that is biased towards \vec{s} in the PPR procedure). In general, teleportation results in a bias to the random walk, which one usually tries to minimize when detecting communities. (See Ref. [125] for clever ways to choose \vec{s} with this goal in mind.)

The algorithm of Refs. [123, 124] deliberately exploits the bias from teleportation to achieve localized results. It computes an approximation to the solution of Eq. (B1) (i.e., it computes an *approximate PPR vector*) by strategically “pushing” mass between the iteratively-updated approximate solution vector and a residual vector in such a way that most of the nodes in the original network are *not* reached. Consequently, this algorithm is typically *much* faster for moderately-large to very large graphs than is the naïve algorithm to compute a solution to Eq. (B1). The algorithm is parametrized in terms of a “truncation” parameter ϵ where larger values of ϵ correspond to more locally-biased solutions. We refer to this procedure as the ACL CUT method.

2. Dynamics Type 2: Local Spectral Partitioning (the “MOV CUT” method).

In this procedure, we formalize the idea of a locally-biased version of the leading nontrivial eigenvector of the normalized Laplacian \mathcal{L} that can be used in a locally-biased version of traditional spectral graph partitioning.

Following Ref. [37], consider the following optimization problem:

$$\begin{aligned} & \underset{\vec{x}}{\text{minimize}} && \vec{x}^T \mathcal{L} \vec{x} \\ & \text{subject to} && \vec{x}^T \vec{x} = 1, \\ & && \vec{x}^T D^{1/2} \vec{1} = 0 \\ & && (\vec{x}^T D^{1/2} \vec{s})^2 \geq \kappa, \end{aligned} \quad (\text{B2})$$

where κ is a locality parameter and \vec{s} is a vector, which satisfies the constraints $\vec{s}^T D \vec{s} = 1$ and $\vec{s}^T D \vec{1} = 0$, and which represents a seed set of nodes. That is, in the norm defined by the diagonal D matrix, the seed vector \vec{s} is unit length and is exactly orthogonal to the all-ones vector. This *locally-biased* version of standard spectral graph partitioning (which becomes the usual global spectral-partitioning problem if the locality constraint $(\vec{x}^T D^{1/2} \vec{s})^2 \geq \kappa$ is removed) was introduced in [37], where it was shown that the solution vector \vec{x}^* inherits many of the nice properties of the solution to the usual global spectral-partitioning problem. The solution \vec{x}^* is of the form

$$\vec{x}^* = c(L_G - \gamma D_G)^+ D_G \vec{s}, \quad (\text{B3})$$

where the parameter $\gamma \in (-\infty, \lambda_2(G))$ is related to the teleportation parameter α via the relation $\gamma = \frac{\alpha-1}{\alpha}$ (see [37]) and $c \in [0, \infty]$ is a normalization constant.

As one can see from Eq. (B3), the solution \vec{x}^* of Eq. (B2) is an *exact PPR vector* with personalized teleportation vector \vec{s} . Consequently, it can be computed as the solution to a system of linear equations. In addition, if one performs a sweep cut (see the discussion below) of this solution vector to obtain a locally-biased network partition, then one obtains Cheeger-like guarantees on approximation quality for the associated network community. Moreover, if the seed vector \vec{s} corresponds to the indicator vector of a single node i , then this is a relaxation of the following *locally-biased graph partitioning problem*: given as input a graph $G = (V, E, w)$, an input node u , and a positive integer k ; find a set of nodes $S \subseteq V$ that is the best conductance set of nodes of volume no greater than k that contains the input node i [37]. We refer to this procedure (with a seed vector corresponding to a single seed node) as the MOV CUT method.

3. Dynamics Type 3: Local Geodesic Spreading (the “EGO NET” method).

In this procedure, we perform a geodesic-based (i.e., ego-network-based) dynamics that is analogous to the local random walks that we described above. This method is similar to the technique for finding local communities that was introduced in Ref. [38] and that was generalized to weighted networks in Ref. [126]. Starting with a seed node s and a distance parameter k , this method considers all nodes j whose geodesic distance from s is at most k away—i.e., all nodes j such that $\Delta_{sj} \leq k$ —to form a local community. In the unweighted case, the *egocentric network* (i.e., *ego network* or *ego-net* [127]) for a seed node (the ego) is the subgraph induced by the seed node’s 1-neighborhood—i.e., the network that consists of all nodes that are in the 1-neighborhood (including the seed node) and all edges between these nodes that are present in the original network. (The traditional definition of an ego-net excludes the seed node and its edges, but we specifically include them.) We use the term *k-ego-net* for the

subgraph that is induced by the k -neighborhood of a seed node. Consequently, the local communities that we obtain using this method are simply the k -ego-nets of the seed node. For consistency with the other two methods, it is useful to think of this method as inducing a ranking of the nodes:

$$\text{EgoRank}_i(s) = \frac{1}{1 + \Delta_{is}}, \quad (\text{B4})$$

where i is some node in a network. Given the ranking interpretation in Eq. (B4), we recover local geodesic-based communities from the EgoRank vector by using the sweep cut procedure that we describe below. The underlying dynamics for this method is analogous to the extreme case of a susceptible-infected (SI) spreading process [43, 128], in which an infected node infects all its neighbors with probability 1 at the time step following the one in which it is infected. One can then interpret the EgoRank of node i for a seed node s as the inverse of the time that it takes for node i to first become infected when only the seed node s is infected initially. We refer to this procedure as the EGONET method.

4. Sampling Procedures and Parameter Choices

To obtain an accurate picture of local community structure at different size scales throughout a network, we run each of the above community-identification procedures many times, starting at different seed nodes and running for different numbers of steps, and we then examine which nodes get visited as the dynamical processes unfold. For each seed node and value of the parameters, each of the ACLCUT, MOV CUT, and EGONET methods returns a vector that can be used to “rank” the nodes of a network (in a locally-biased and size-resolved manner): ACLCUT and MOV CUT return a variant of the PPR vector, and EGONET returns the EgoRank vector in Eq. (B4). Then, given a ranking vector \vec{p} , the so-called “sweep sets” are given by $S_t = \{i \in V : \vec{p}_i \geq t\}$; and thus there are at most $n + 1$ distinct sweep sets (where we recall that n is the number of nodes in the graph). A corresponding “sweep cut” is then the partition of the network obtained from a sweep set that has minimal conductance, over all $n + 1$ possible sweep set partitions. By computing the conductance for each of the sweep sets, one obtains a locally-biased estimate for an NCP, centered around a seed node. One can then estimate a global NCP by taking the lower envelope over local NCPs for different seed nodes and parameter values. Our MATLAB code that implements these methods is available at [129].

Recall that ACLCUT has two parameters (the teleportation parameter α and the truncation parameter ϵ), but that MOV CUT only has a single parameter (a teleportation parameter).

For ACLCUT, theoretical results [123] suggest that the method should find good communities of volume roughly ϵ^{-1} , where we have ignored constants and logarithmic

factors. Furthermore, for a seed node i with strength k_i , ACLCUT returns empty communities for $\epsilon < k_i^{-1}$. This suggests that sampling using $\epsilon \in [k_{\max}^{-1}, \text{vol}(G)^{-1}]$ gives good coverage of different size scales in practice. In this paper, we use 20 logarithmically-spaced points in $[k_{\max}^{-1}, \text{vol}(G)^{-1}]$ (including the endpoints) to generate Figs. 6, 7, and 15. In addition, we use $\tilde{\alpha} = 0.001$, where $\tilde{\alpha}$ is the teleportation parameter of the “lazy random walk” defined in [123]. The (conventional) teleportation parameter that we use satisfies $\alpha = 1 - \frac{2\tilde{\alpha}}{1+\tilde{\alpha}}$, so that $\alpha \approx 0.998$ in Eq. (B1). In our computations, we observed that increasing α leads to more accurate NCPs at the cost of longer computation times.

For MOV CUT, we use 20 equally-spaced values of α in the interval $[0.7, (1 - \lambda_2)^{-1} - 10^{-10}]$ (including the endpoints), where $(1 - \lambda_2)^{-1}$ is the theoretical maximum for α (see [37]).

To sample seed nodes, we modified the strategy described in Ref. [25] to be applicable to the MOV CUT method as well as the ACLCUT method. For each choice of parameter values, we sampled nodes uniformly at random without replacement and stopped the sampling process either when all nodes were sampled or when the sampled local communities sufficiently covered the entire network. To determine sufficient coverage, we tracked how many times each node was included in the best local community that we obtained from the sweep sets and stopped the procedure once each node was included at least 10 times. This procedure ensures that good communities are sampled consistently.

The EGONET method does not have any size-scale parameters. For the network sizes that we consider, it is feasible to use all nodes rather than sampling them. We use this approach to generate Figs. 18 and 19.

Finally, for readability, we only plotted the NCPs for communities that contain at most half of the nodes in a network. The symmetry in the definition of conductance (see Eq. (3)) implies that the complement of a good small community is necessarily a good large community and vice versa. Hence, a sampled NCP is roughly symmetric, though this is hard to see on a logarithmic scale, and an NCP without sampling is necessarily symmetric.

Appendix C: Detailed Results for the MOV CUT Method

The MOV CUT method provides an alternative way of sampling local community profiles to construct an NCP. Unlike ACLCUT, which uses *only* local information to obtain good communities, MOV CUT also incorporates some global information about a network to construct local communities around a seed node. In particular, this implies that there can be sweep sets and thus communities that consist of disconnected components of a network. Such communities have infinitely large conductance ratios. We observe this phenomenon often for the coauthorship and Facebook networks, but it almost never




FIG. 16. NCP plots [in panel (a)] and CRP plots [in panel (b)] for FB-JOHNS55, CA-GRQC, and US-SENATE (i.e., the smaller network in each of the three pairs of networks from Table I) generated using the MOV CUT method. The thin curves are the NCPs that we obtain when also consider disconnected sweep sets.

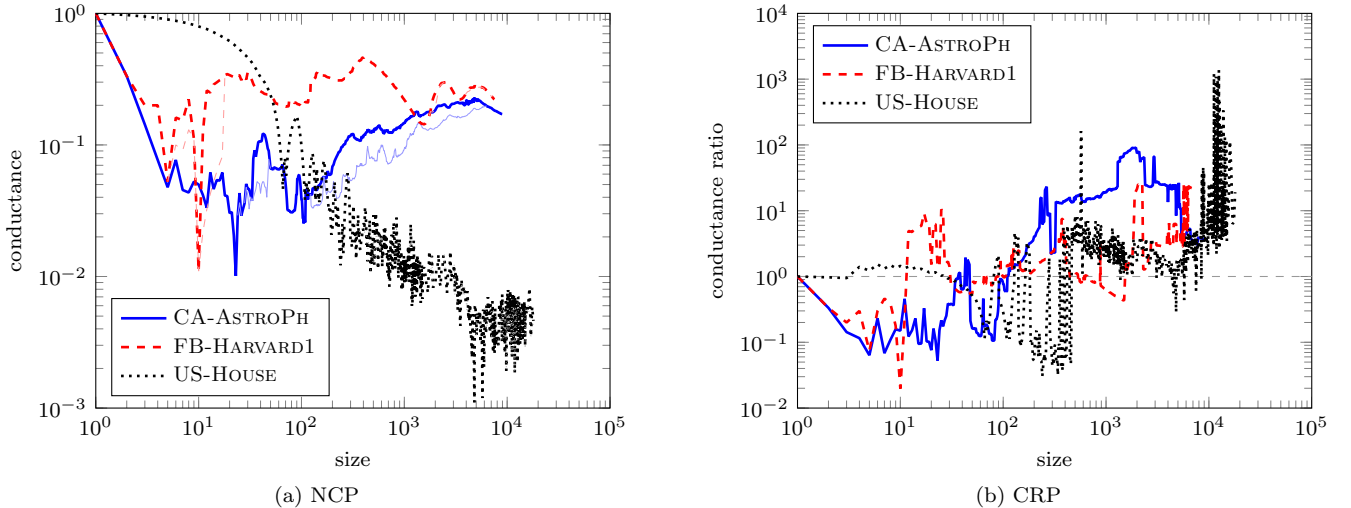


FIG. 17. NCP plots [in panel (a)] and CRP plots [in panel (b)] for CA-ASTROPH, FB-HARVARD1, and US-HOUSE (i.e., the larger network in each of the three pairs of networks from Table I) generated using the MOV CUT method. The thin curves are the NCPs that we obtain when also consider disconnected sweep sets.

occurs for the Congressional voting networks. Upon examination, these sweep sets consist of several small sets of peripheral nodes, each of which has moderate to very low conductance, but which are otherwise unrelated. Although one would not usually think of such a set of nodes as a single good community, optimization-based algorithms often clump several unrelated communities into a single community for networks with a global core-periphery structure. For completeness and comparison, we include our results both when we keep the disconnected sweep sets and when we restrict our attention to connected communities. As we discuss below, the NCP does not change substantially, although there are some

small differences.

The resulting NCPs for the MOV CUT method (see Figs. 16a and 17a) are similar to those that we obtained for the ACL CUT method (see Figs. 6a and 7a), although there are a few differences worth discussing. The CRP plots are also very similar (compare Figs. 16b and 17b to Figs. 6b and 7b). For the coauthorship networks (CA-GRQC and CA-ASTROPH), as well as FB-HARVARD1, both MOV CUT and ACL CUT identify the same good small communities that are responsible for the spikes in the NCP plots. In addition, the communities that yield the dips in the NCPs for FB-JOHNS55 near 220 and 1100 nodes, and for FB-HARVARD1 near 1500 nodes, all share




FIG. 18. NCP plots [in panel (a)] and CRP plots [in panel (b)] for CA-GRQC, FB-JOHN55, and US-SENATE (i.e., the smaller network in each of the three pairs of networks from Table I) using the EGONET method. We find qualitatively similar behavior as with the other two methods, although the NCPs are shifted upwards and some of the large-scale structure is no longer present (especially in the Facebook network).




FIG. 19. NCP plots [in panel (a)] and CRP plots [in panel (b)] for CA-ASTROPH, FB-HARVARD1, and US-HOUSE (i.e., the larger networks in each of the three pairs of networks from Table I) using the EGONET method. We find qualitatively similar behavior as with the other two methods, although the NCPs are shifted upwards and some of the large-scale structure is no longer present (especially in the Facebook network).

more than 98% of their nodes. This indicates that both methods are able to find roughly the same community-like structures. However, the results from the MOV-CUT NCP for CA-GRQC is higher and less choppy than the one that we computed using ACLCUT—because the truncation employed by ACLCUT performs a form of implicit sparsity-based regularization that is absent from MOV-CUT. See Refs. [101, 102, 130] for a discussion and precise characterization of this regularization. For the coauthorship and Facebook networks, we also note that there are regions of the computed NCPs, when using the MOV-

CUT method, in which one finds disconnected sweep sets (see the thin curves) with lower conductance than that for the best connected sets of the same size. At other sizes, we see some differences between the NCPs from MOV-CUT and ACLCUT. This illustrates that the two methods can have somewhat different local behavior, although both methods produce similar insights regarding the large-scale structure in these networks. In Section IV C, we discuss some of these differences between our results from the two methods in more detail.

Appendix D: Detailed Results for the EGoNET Method

The EGoNET method was not originally developed to optimize conductance, although there is some recent evidence that k -neighborhoods can be good conductance communities [131]. The assumption that underlies the EGoNET method is that nodes in the same community should be connected by short paths. However, unlike the spectral-based methods (ACLcUT and MoVcUT), the EGoNET method does not take into account the number of paths between nodes. In contrast to Ref. [131], which considered only 1-neighborhoods, here we also examine k -neighborhoods with $k > 1$. We can then use this method to sample a complete NCP for a network.

Despite its simplicity, and in agreement with Ref. [131], the EGoNET method produces NCP's that are qualitatively similar to those from both the ACLcUT and MoVcUT methods, for all of the networks that we considered; see Figs. 18 and 19. The NCPs for the EGoNET method are shifted upwards compared to those for the ACLcUT and MoVcUT methods; and this is particularly noticeable at larger community size. This is unsurprising, because the latter two methods more aggressively optimize the conductance objective. However, for all six of our networks, this method preserves an NCP's small-scale structure as well as the global tendency to be upward-sloping, flat, or downward-sloping. This provides further evidence that the qualitative features of an NCP pro-

vide a signature of community structure in a network and are not just an artifact of a particular way to sample communities. In Section IV C, we give a more detailed comparison between the results of these methods.

ACKNOWLEDGEMENTS

LGSJ acknowledges a CASE studentship award from the EPSRC (BK/10/039). MAP was supported by a research award (#220020177) from the James S. McDonnell Foundation, the EPSRC (EP/J001759/1), and the FET-Proactive project PLEXMATH (FP7-ICT-2011-8; grant #317614) funded by the European Commission; MAP also thanks SAMSI for supporting several visits and MWM for his hospitality during his sabbatical at Stanford. PJM was funded by the NSF (DMS-0645369) and by Award Number R21GM099493 from the National Institute Of General Medical Sciences. MWM acknowledges funding from the Army Research Office and from the Defense Advanced Research Projects Agency. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies. In addition, we thank Adam D'Angelo and Facebook for providing the Facebook data, Keith Poole for providing the Congressional voting data (which is available from Ref. [132]), and Jure Leskovec for making many large network data sets publicly available as part of SNAP [85].

-
- [1] M. O. Jackson, *Social and Economic Networks* (Princeton University Press, Princeton, 2008).
 - [2] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World* (Cambridge University Press, Cambridge, 2010).
 - [3] M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).
 - [4] A.-L. Barabási, *Nat. Phys.* **1**, 68 (2005).
 - [5] M. A. Porter, J.-P. Onnela, and P. J. Mucha, *Notices Amer. Math. Soc.* **56**, 1082 (2009).
 - [6] S. Fortunato, *Phys. Rep.* **486**, 75 (2010).
 - [7] M. Girvan and M. E. J. Newman, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 7821 (2002).
 - [8] M. A. Porter, P. J. Mucha, M. E. J. Newman, and C. M. Warmbrand, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 7057 (2005).
 - [9] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, *Science* **328**, 876 (2010).
 - [10] K. T. Macon, P. J. Mucha, and M. A. Porter, *Phys. A* **391**, 343 (2012).
 - [11] A. L. Traud, P. J. Mucha, and M. A. Porter, *Phys. A* **391**, 4165 (2012).
 - [12] A. L. Traud, E. D. Kelsic, P. J. Mucha, and M. A. Porter, *SIAM Rev.* **53**, 526 (2011).
 - [13] M. C. González, H. J. Herrmann, J. Kertész, and T. Vicsek, *Phys. A* **379**, 307 (2007).
 - [14] A. C. F. Lewis, N. S. Jones, M. A. Porter, and C. M. Deane, *BMC Syst. Biol.* **4**, 100 (2010).
 - [15] D. S. Bassett, E. T. Owens, K. E. Daniels, and M. A. Porter, *Phys. Rev. E* **86**, 041306 (2012).
 - [16] P. Ronhovde, S. Chakrabarty, D. Hu, M. Sahu, K. K. Sahu, K. F. Kelton, N. A. Mauro, and Z. Nussinov, *Eur. Phys. J. E* **34**, 105 (2011).
 - [17] D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 7641 (2011).
 - [18] N. F. Wymbs, D. S. Bassett, P. J. Mucha, M. A. Porter, and S. T. Grafton, *Neuron* **74**, 936 (2012).
 - [19] D. S. Bassett, N. F. Wymbs, M. P. Rombach, M. A. Porter, P. J. Mucha, and S. T. Grafton, *PLoS Comput. Biol.* **9**, e1003171 (2013).
 - [20] T. S. Evans, R. Lambiotte, and P. Panzarasa, *Scientometrics* **89**, 381 (2011).
 - [21] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7332 (2007).
 - [22] P. Expert, T. S. Evans, V. D. Blondel, and R. Lambiotte, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 7663 (2011).
 - [23] D. Grady, R. Brune, C. Thiemann, F. Theis, and D. Brockmann, in *Handbook of Optimization in Complex Networks*, Springer Optimization and its Applications, edited by M. T. Thai and P. M. Pardalos (Springer, New York, 2012) pp. 169–208.
 - [24] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, in *Proceedings of the 17th International Conference on World Wide Web, WWW '08* (ACM, New York,

- 2008) pp. 695–704.
- [25] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, *Internet Mathematics* **6**, 29 (2009).
- [26] J. Leskovec, K. J. Lang, and M. W. Mahoney, in *Proceedings of the 19th International Conference on World Wide Web*, WWW '10 (ACM, New York, 2010) pp. 631–640.
- [27] We do not wish to tie ourselves too closely to particular numbers, because such numbers can vary dramatically in different classes of networks. As a rule of thumb, however, we think of “small,” “medium-sized,” and “large” communities as those that contain, respectively, tens to about a hundred, hundreds to thousands, and thousands or more nodes. Similarly, we think of “small” networks as having tens to hundreds of nodes and “large” networks as having thousands, tens of thousands, or even more nodes.
- [28] A. Lancichinetti and S. Fortunato, *Phys. Rev. E* **80**, 056117 (2009).
- [29] B. H. Good, Y.-A. de Montjoye, and A. Clauset, *Phys. Rev. E* **81**, 046106 (2010).
- [30] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, *J. Stat. Mech. Theor. Exp.* **10**, P10008 (2008).
- [31] J. Yang and J. Leskovec, in *Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics*, MDS '12 (ACM, New York, 2012) pp. 3:1–3:8.
- [32] R. Lambiotte, J.-C. Delvenne, and M. Barahona, (2009), arXiv:0812.1770 [physics.soc-ph].
- [33] M. T. Schaub, R. Lambiotte, and M. Barahona, *Phys. Rev. E* **86**, 026112 (2012).
- [34] M. T. Schaub, J. C. Delvenne, S. N. Yaliraki, and M. Barahona, *PLoS One* **7**, e32210 (2012).
- [35] S. Fortunato and M. Barthelemy, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 36 (2007).
- [36] J. Yang and J. Leskovec, (2012), arXiv:1205.6228 [cs.SI].
- [37] M. W. Mahoney, L. Orecchia, and N. K. Vishnoi, *J. Mach. Learn. Res.* **13**, 2339 (2012).
- [38] J. Bagrow and E. Bollt, *Phys. Rev. E* **72**, 046108 (2005).
- [39] N. A. Christakis and J. H. Fowler, *N. Engl. J. Med.* **357**, 370 (2007).
- [40] N. A. Christakis and J. H. Fowler, *Statistics in Medicine* **32**, 556 (2013).
- [41] F. Brauer and C. Castillo-Chavez, *Mathematical Models in Population Biology and Epidemiology*, 2nd ed., Texts in Applied Mathematics, Vol. 40 (Springer, New York, 2012).
- [42] R. Lambiotte and M. Rosvall, *Papers in Physics* **5**, 050003 (2013).
- [43] J. P. Gleeson, *Phys. Rev. X* **3**, 021004 (2013).
- [44] J. Ugander, B. Karrer, L. Backstrom, and C. Marlow, (2011), arXiv:1111.4503 [cs.SI].
- [45] M. Gjoka, E. Smith, and C. Butts, (2013), arXiv:1308.3297 [cs.SI].
- [46] T. J. Hansen and M. W. Mahoney, in *Advances in Neural Information Processing Systems*, Vol. 25, edited by P. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (NIPS, 2012) pp. 2537–2545.
- [47] M. Jerrum and A. Sinclair, in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, STOC '88 (ACM, New York, 1988) pp. 235–244.
- [48] K. Lerman and R. Ghosh, *Phys. Rev. E* **86**, 026108 (2012).
- [49] R. Ghosh and K. Lerman, *Discrete and Continuous Dynamical Systems Series B* **19**, 1355 (2014), arXiv:1209.4616 [cs.SI].
- [50] R. Ghosh, K. Lerman, S.-H. Teng, and X. Yan, (2014), arXiv:1406.3387 [cs.SI].
- [51] J. P. Bagrow, *J. Stat. Mech. Theor. Exp.* **2008**, P05001 (2008).
- [52] A. Lancichinetti, F. Radicchi, J. J. Ramasco, and S. Fortunato, *PLoS One* **6**, e18961 (2011).
- [53] M. Rosvall and C. T. Bergstrom, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 1118 (2008).
- [54] R. Lambiotte, R. Sinatra, J.-C. Delvenne, T. S. Evans, M. Barahona, and V. Latora, *Phys. Rev. E* **84**, 017102 (2011).
- [55] E. Ziv, M. Middendorf, and C. H. Wiggins, *Phys. Rev. E* **71**, 046117 (2005).
- [56] C. Borgs, M. Brautbar, J. Chayes, S. Khanna, and B. Lucier, in *Internet and Network Economics*, Lecture Notes in Computer Science, Vol. 7695 (Springer, Berlin Heidelberg, 2012) pp. 406–419.
- [57] S. Melnik, A. Hackett, M. A. Porter, P. J. Mucha, and J. P. Gleeson, *Phys. Rev. E* **83**, 036112 (2011).
- [58] M. R. Bridson and A. Häflicher, *Metric Spaces of Non-Positive Curvature* (Springer, Berlin, 1999).
- [59] W. W. Zachary, *J. Anthropol. Res.* **33**, 452 (1977).
- [60] W. Chen, W. Fang, G. Hu, and M. W. Mahoney, *Internet Mathematics* **9**, 434 (2013).
- [61] The parameter choices for the block models have subtle effects that are not explained by the relative sizes of the probabilities α_{11} , α_{12} , and α_{22} . For a network with pronounced core-periphery structure (such as the one we show in Fig. 1b), where $\alpha_{11} \gg \alpha_{12} \gg \alpha_{22}$, one additionally needs $n_1\alpha_{11} \gg 1$ (for the core to be sufficiently connected), $n_1\alpha_{12} \ll 1$ (otherwise most peripheral nodes are directly connected to the core), and $n_2\alpha_{22} \approx 1$ (such that the periphery is reasonably connected but not too dense). Other parameter choices that respect the basic constraint on the probabilities still produce networks with well-defined density-based core-periphery structure [62–64], but they do not “look like” Fig. 1b (they tend to resemble Fig. 1c, provided they are reasonably well-connected). For an expander-like network (such as the one we show in Fig. 1c), where $\alpha_{11} \approx \alpha_{12} \approx \alpha_{22} \approx \alpha$, one needs $(n_1 + n_2)\alpha \gg 1$ for the network to be sufficiently well-connected.
- [62] P. Csermely, A. London, L.-Y. Wu, and B. Uzzi, *J. Complex Networks* **1**, 93 (2013).
- [63] S. P. Borgatti and M. G. Everett, *Soc. Networks* **21**, 375 (2000).
- [64] M. P. Rombach, M. A. Porter, J. H. Fowler, and P. J. Mucha, *SIAM J. App. Math.* **74**, 167 (2014).
- [65] M. E. J. Newman, *Phys. Rev. E* **74**, 036104 (2006).
- [66] M. P. Rombach, “The Zachary karate cake,” (2012), available at <http://people.maths.ox.ac.uk/porterm/temp/karatecake.jpg>.
- [67] C. R. Shalizi, A. Hagberg, and A. Clauset, “Network scientists with karate trophies,” (2013), available at <http://networkkarate.tumblr.com>.
- [68] H. A. Simon, *P. Am. Philos. Soc.* **106**, 467 (1962).
- [69] F. R. K. Chung, *Spectral graph theory*, CBMS Regional Conference Series in Mathematics, Vol. 92 (AMS, Providence, 1997).

- [70] F. R. K. Chung and L. Lu, *Complex graphs and networks*, CBMS Regional Conference Series in Mathematics, Vol. 107 (AMS, Providence, 2006).
- [71] S. Arora, S. Rao, and U. Vazirani, *J. ACM* **56**, 5:1 (2009).
- [72] D. J. Watts, *Am. J. Sociol.* **105**, 493 (1999).
- [73] A. Condon and R. M. Karp, *Random Struct. Alg.* **18**, 116 (2001).
- [74] M. Mihail, in *Proceedings of the 30th Annual Symposium on Foundations of Computer Science* (IEEE, 1989) pp. 526–531.
- [75] Z. A. Zhu, S. Lattanzi, and V. Mirrokni, in *Proceedings of the 30th International Conference on Machine Learning*, edited by S. Dasgupta and D. McAllester (JMLR W&CP, 2013) pp. 396–404.
- [76] P. J. Mucha and M. A. Porter, *Chaos* **20**, 041108 (2010).
- [77] M. Kivela, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, *J. Complex Networks* **2**, 203 (2014).
- [78] D. S. Bassett, M. A. Porter, N. F. Wymbs, S. T. Grafton, J. M. Carlson, and P. J. Mucha, *Chaos* **23**, 013142 (2013).
- [79] M. De Domenico, A. Solé-Ribalta, S. Gómez, and A. Arenas, *Proc. Natl. Acad. Sci. U.S.A.* **111**, 8351 (2014).
- [80] M. De Domenico, A. Solé-Ribalta, E. Cozzo, M. Kivela, Y. Moreno, M. A. Porter, S. Gómez, and A. Arenas, *Phys. Rev. X* **3**, 041022 (2013).
- [81] S. Gómez, A. Diaz-Guilera, J. Gómez-Gardenes, C. J. Pérez-Vicente, Y. Moreno, and A. Arenas, *Phys. Rev. Lett.* **110**, 028701 (2013).
- [82] J.-P. Onnela, J. Saramäki, J. Kertész, and K. Kaski, *Phys. Rev. E* **71**, 065103 (2005).
- [83] J. Saramäki, M. Kivela, J.-P. Onnela, K. Kaski, and J. Kertész, *Phys. Rev. E* **75**, 027105 (2007).
- [84] D. J. Watts and S. H. Strogatz, *Nature* **393**, 440 (1998).
- [85] J. Leskovec, “Stanford network analysis project,” Available at <http://snap.stanford.edu>.
- [86] I. S. Jutla, L. G. S. Jeub, and P. J. Mucha, “A generalized Louvain method for community detection implemented in MATLAB,” (2012), available at <http://netwiki.amath.unc.edu/GenLouvain>.
- [87] Banded matrices arise, for example, in the study of finite-difference equations. In banded matrices, the non-zero entries are confined to diagonal bands near the main diagonal. Using our informal terminology from Section II B, this is an extreme example of a “hot dog”.
- [88] T. Kamada and S. Kawai, *Inform. Process. Lett.* **31**, 7 (1989).
- [89] L. G. S. Jeub, “Spring-based visualization for networks with communities,” (2014), available at <https://github.com/LJeub/SpringVisCom>.
- [90] The two small “communities” in FB-HARVARD1 are responsible for the low value of λ_2 for this network compared with all of the other Facebook networks (recall Fig. 4). Removing these two communities from the network increases the value of λ_2 by an order of magnitude (from 0.0094 to 0.075), yielding a value that is fairly typical for networks in the FACEBOOK100 data set. In other words, aside from the 15 nodes from these two “communities,” FB-HARVARD1 looks like FB-JOHNS55 and the rest of the FACEBOOK100 networks. This is one of few examples where a feature of an NCP is *not* robust.
- [91] For FB-JOHNS55, the community of size about 220 corresponds closely to a set of students with the same major, and the community associated with the dip near 1100 corresponds to first-year students. Similarly, the community associated with the dip near 1500 for FB-HARVARD1 also corresponds to first-year students. Consistent with the results of Ref. [11], we find similar poorly-connected and moderately-large communities that correlate reasonably well with first-year students in most of the other networks in the FACEBOOK100 data set. (As usual, the most notable exception among the FACEBOOK100 networks is CALTECH36, which is known to be influenced much more by dormitory (“House”) residence than by class year [11].)
- [92] M. Cucuringu and M. W. Mahoney, (2011), arXiv:1109.1355 [cs.DM].
- [93] C. Spearman, *Am. J. Psychol.* **15**, 72 (1904).
- [94] L. Backstrom, P. Boldi, M. Rosa, J. Ugander, and S. Vigna, in *Proceedings of the 3rd Annual ACM Web Science Conference*, WebSci ’12 (ACM, New York, 2012) pp. 33–42.
- [95] This holds for all seed nodes because the maximum voting similarity between a pair of Senators is 1 (i.e., they voted the same way on every bill) and all interlayer edges (i.e., edges between the same Senator in different Congresses) have a weight of 1. As one reduces the weight of interlayer edges, one needs to choose an increasingly large value of k for nodes from interlayer edges to appear in the k -neighborhood. Hence, an increasing number of Senators from the same Congress can appear in a k -neighborhood that does not contain any nodes from interlayer edges.
- [96] M. Sales-Pardo, R. Guimerà, A. A. Moreira, and L. A. N. Amaral, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 15224 (2007).
- [97] J. Reichardt and S. Bornholdt, *Phys. Rev. E* **74**, 016110 (2006).
- [98] A. Arenas, A. Fernández, and S. Gomez, *New J. Phys.* **10**, 053039 (2008).
- [99] Z. Bar-Joseph, D. K. Gifford, and T. S. Jaakkola, *Bioinformatics* **17**, S22 (2001).
- [100] A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter, (2009), arXiv:0907.3509 [physics.soc-ph].
- [101] M. W. Mahoney and L. Orecchia, in *Proceedings of the 28th International Conference on Machine Learning*, ICML ’11, edited by L. Getoor and T. Scheffer (ACM, New York, 2011) pp. 121–128.
- [102] D. Gleich and M. W. Mahoney, in *Proceedings of the 31st International Conference on Machine Learning*, edited by T. Jebara and E. P. Xing (JMLR W&CP, 2014) pp. 1018–1025.
- [103] Y.-Y. Ahn, J. P. Bagrow, and S. Lehmann, *Nature* **466**, 761 (2010).
- [104] I. Psorakis, S. Roberts, M. Ebdon, and B. Sheldon, *Phys. Rev. E* **83**, 066114 (2011).
- [105] B. Ball, B. Karrer, and M. E. J. Newman, *Phys. Rev. E* **84**, 036103 (2011).
- [106] S. Lehmann, “Intermittent updates: Pervasive overlap,” (2010), available at <http://sunelehmann.com/2010/06/29/pervasive-overlap/>.
- [107] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani, *J. Mach. Learn. Res.* **11**, 985 (2010).

- [108] A. Lancichinetti, S. Fortunato, and F. Radicchi, *Phys. Rev. E* **78**, 046110 (2008).
- [109] A. Lancichinetti and S. Fortunato, *Sci. Rep.* **2**, 336 (2012).
- [110] A. Lancichinetti and S. Fortunato, *Phys. Rev. E* **80**, 016118 (2009).
- [111] A. Lancichinetti, “Benchmarks (binary networks v.9),” (2010), available at <http://sites.google.com/site/andrealancichinetti/software>.
- [112] To give one example, the coauthorship networks that we examined possess a familiar upward-sloping NCP and a nested core-periphery structure with small, well-separated communities in the periphery. However, the small and large communities that we observed in the Facebook networks do not lead to bottlenecks for diffusion processes. This novel NCP behavior in the Facebook graphs is consistent both with the larger edge density of the Facebook networks relative to the SNAP graphs [24–26] as well as with previous work on simple dynamical processes on these Facebook networks [57].
- [113] G. K. Orman, V. Labatut, and H. Cherifi, *Int. J. Web Based Communities* **3**, 349 (2013).
- [114] M. S. Granovetter, *Am. J. Sociol.* **78**, 1360 (1973).
- [115] P. Sarnak, *Not. Amer. Math. Soc.* **51**, 762 (2004).
- [116] S. Hoory, N. Linial, and A. Wigderson, *Bull. Amer. Math. Soc.* **43**, 439 (2006).
- [117] A. Arenas, A. Díaz-Guilera, and C. J. Pérez-Vicente, *Phys. Rev. Lett.* **96**, 114102 (2006).
- [118] R. Lambiotte, in *Proceedings of the 8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, WiOpt (IEEE, 2010)* pp. 546–553.
- [119] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 12755 (2010).
- [120] E. Estrada, *Chaos* **21**, 016103 (2011).
- [121] D. Anderson, A. Tenzer, G. Barlev, M. Girvan, T. M. Antonsen, and E. Ott, *Chaos* **22**, 013102 (2012).
- [122] M. Beguerisse-Díaz, G. Garduño-Hernández, B. Vangelov, S. N. Yaliraki, and M. Barahona, (2013), arXiv:1311.6785 [physics.soc-ph].
- [123] R. Andersen, F. Chung, and K. J. Lang, in *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, FOCS '06 (IEEE, 2006)* pp. 475–486.
- [124] R. Andersen and K. J. Lang, in *Proceedings of the 15th International Conference on World Wide Web, WWW '06 (ACM, New York, 2006)* pp. 223–232.
- [125] R. Lambiotte and M. Rosvall, *Phys. Rev. E* **85**, 056107 (2012).
- [126] M. A. Porter, P. J. Mucha, M. E. J. Newman, and A. J. Friend, *Phys. A* **386**, 414 (2007).
- [127] S. Wasserman and K. Faust, *Social Network Analysis: Methods and Applications* (Cambridge University Press, Cambridge, 1994).
- [128] M. Karsai, M. Kivela, R. K. Pan, K. Kaski, J. Kerész, A.-L. Barabási, and J. Saramäki, *Phys. Rev. E* **83**, 025102 (2011).
- [129] L. G. S. Jeub, “LocalCommunities,” (2014), available at <https://github.com/LJeub/LocalCommunities>.
- [130] M. W. Mahoney, in *Proceedings of the 31st ACM Symposium on Principles of Database Systems, PODS '12 (ACM, New York, 2012)* pp. 143–154.
- [131] D. F. Gleich and C. Seshadhri, in *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12 (ACM, New York, 2012)* pp. 597–605.
- [132] K. T. Poole, “Voteview,” (2014), available at <http://voteview.com>.