

# *A local perspective on community structure in multilayer networks*

LUCAS G. S. JEUB

*Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute,  
University of Oxford, OX2 6GG, UK  
and  
Center for Complex Networks and Systems Research, School of Informatics  
and Computing, Indiana University, Bloomington, IN 47408, USA  
(e-mail: lucasjeub@gmail.com)*

MICHAEL W. MAHONEY

*International Computer Science Institute, Berkeley, CA 94704, USA  
and  
Department of Statistics, University of California at Berkeley, Berkeley, CA 94720, USA  
(e-mail: mahoneymw@gmail.com)*

PETER J. MUCHA

*Carolina Center for Interdisciplinary Applied Mathematics, Department of Mathematics, University of  
North Carolina, Chapel Hill, NC 27599-3250, USA  
(e-mail: mucha@unc.edu)*

MASON A. PORTER

*Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute,  
University of Oxford, OX2 6GG, UK,  
CABDyN Complexity Centre, University of Oxford, Oxford, OX1 1HP, UK  
and  
Department of Mathematics, University of California, Los Angeles, Los Angeles, California 90095, USA  
(e-mail: mason@math.ucla.edu)*

---

## Abstract

The analysis of multilayer networks is among the most active areas of network science, and there are several methods to detect dense “communities” of nodes in multilayer networks. One way to define a community is as a set of nodes that trap a diffusion-like dynamical process (usually a random walk) for a long time. In this view, communities are sets of nodes that create bottlenecks to the spreading of a dynamical process on a network. We analyze the local behavior of different random walks on multiplex networks (which are multilayer networks in which different layers correspond to different types of edges) and show that they have very different bottlenecks, which correspond to rather different notions of what it means for a set of nodes to be a good community. This has direct implications for the behavior of community-detection methods that are based on these random walks.

**Keywords:** community structure, clustering, multilayer networks, random walks, network community profiles, spectral graph theory, personalized PageRank

---

A “community” in a network describes a densely connected set of nodes (often relative to a null model), and communities can reveal regularities in processes of network formation, strongly influence the behavior of dynamical processes that take

place on a network, and can correlate with or reveal functional groups of nodes (Porter et al., 2009; Fortunato, 2010; Coscia et al., 2011; Fortunato & Hric, 2016). One can examine community structure in a network from either a global perspective or a local one. When using a global perspective, one typically partitions a network into a set of (potentially overlapping) communities; by contrast, when taking a local perspective, one seeks to determine the community or communities associated with a given node (or set of nodes). A local perspective naturally allows the detection of overlapping communities, as local communities for different seed nodes can share nodes without having to be identical.

From either a global or local perspective, communities can be viewed as dependent not only on network structure but also on dynamical processes (as a surrogate for function) on a network. Moreover, the choices of both dynamical process and initial conditions are very important (Jeub et al., 2015). A popular and successful approach for identifying community structure—both globally (Rosvall & Bergstrom, 2008; Delvenne et al., 2010) and locally (Andersen et al., 2006; Leskovec et al., 2009; Jeub et al., 2015)—is to analyze the behavior of a diffusion, random walk, or other spreading process on a network. This exploits the connection between the presence of communities in a network and the behavior of associated dynamical processes on that network (Lambiotte et al., 2009, 2015; Kloumann et al., 2016).

Although numerous tools have been developed for the analysis of networks (Newman, 2010), most of them concentrate on time-independent networks with only a single type of tie between entities. Such ordinary networks are often unable to capture the complex interactions among entities in the real world. In general, interactions (and the entities themselves) can change over time, and there can also be multiple types of interactions between the same pair of entities. Temporal networks allow one to examine the former situation (Holme & Saramäki, 2012; Holme, 2015), and multiplex networks allow consideration of the latter (Wasserman & Faust, 1994). The use of *multilayer networks* (Boccaletti et al., 2014; Kivela et al., 2014) allows one to examine temporal networks and multiplex networks. In the former case, each layer represents a time or a time window (though it is important to think about issues such as discrete versus continuous time). In the latter case, each layer represents a type of interaction. One can also represent a multiplex temporal network by using a multilayer framework.

Because multilayer networks are graphical structures with nodes and edges, the notion of bottlenecks to dynamical processes on networks extends in a natural way to multilayer networks. (See Boccaletti et al., 2014; Kivela et al., 2014; Salehi et al., 2015; De Domenico et al., 2016 for discussions of numerous dynamical processes on multilayer networks.) The notion that diffusion-like dynamics should exhibit bottlenecks when there are good communities has been a fruitful perspective for generalizing algorithmic detection of global community structure from single-layer networks to multilayer networks (Mucha et al., 2010; De Domenico et al., 2015). In the present paper, we view community structure in multilayer networks from a local perspective, and we demonstrate that one can directly apply the methodology from Jeub et al. (2015) by considering a dynamical process that traverses both intralayer and interlayer edges. (One can also study community structure in multilayer networks using other approaches, such as ones based on stochastic block models (Peixoto, 2015).) One way to do this is to define an appropriate random walk on a multilayer network.

In the present paper, we examine some features of community structure that can occur in multiplex networks. As examples, we use two different random walks to explore the structure of synthetic benchmark multiplex networks and two empirical multiplex networks. In Section 1, we discuss random walks on multilayer networks in general and the two random walks that we explore in detail. In Section 2, we introduce our methodology for identifying and summarizing community structure in networks. We then use this methodology to explore the behavior of the different random walks on synthetic benchmark networks in Section 3 and on a transportation and a social multiplex network in Section 4. We conclude in Section 5.

## 1 Random walks on multilayer networks

As with a traditional network, different choices are possible when defining a random walk on a multilayer network (Mucha et al., 2010; De Domenico et al., 2014; De Domenico et al., 2015; Kuncheva & Montana, 2015). Following Kivelä et al. (2014), a *multilayer network*  $M(V_M, E_M, V, L)$  is a graph  $G_M(V_M, E_M)$  with an additional layer structure. Here,  $V$  is a set of *nodes*,  $L$  is a set of *layers*,  $V_M \subseteq V \times L$  is a set of *state nodes*,<sup>1</sup> and  $E_M \subseteq V_M \times V_M$  is a set of (directed) *edges*. We use  $i\alpha \in V_M$  to denote the state node that represents node  $i \in V$  in layer  $\alpha \in L$  and  $(i\alpha, j\beta) \in E_M$  to denote a directed edge from state node  $i\alpha$  to state node  $j\beta$ . One can encode the connectivity structure of a multilayer network, including both intralayer and interlayer edges, using an *adjacency tensor*  $A$  (the analogue of the adjacency matrix for single-layer networks) with elements

$$A_{j\beta}^{i\alpha} = \begin{cases} 1, & (i\alpha, j\beta) \in E_M \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

One can write a discrete-time random walk on a multilayer network as

$$p_{i\alpha}(t+1) = \sum_{j\beta \in V_M} P_{i\alpha}^{j\beta} p_{j\beta}(t), \quad (2)$$

where  $p_{j\beta}(t)$  is the probability for a random walker to be at node  $j$  in layer  $\beta$  at time  $t$  and  $P_{i\alpha}^{j\beta}$  is the probability for a random walker at node  $j$  in layer  $\beta$  to transition to node  $i$  in layer  $\alpha$  in one time step. The *transition tensor*  $P$  encodes both the intralayer and interlayer behavior of a random walk. We also want the random walk to be ergodic, so that it has a well-defined stationary distribution  $p_{i\alpha}(\infty)$ . The stationary distribution is a fixed point of Equation (2). That is, it satisfies

$$p_{i\alpha}(\infty) = \sum_{j\beta \in V_M} P_{i\alpha}^{j\beta} p_{j\beta}(\infty). \quad (3)$$

There are different ways to define a random walk on a multilayer network that reduce to the usual definition of a random walk for a single-layer network. The most direct way to generalize the concept of a random walk to a multilayer network is the *classical random walk* (Mucha et al., 2010; De Domenico et al., 2014), which treats interlayer edges and intralayer edges as equivalent objects (though they can be

<sup>1</sup> Following De Domenico et al. (2015), we use the term *state node* to refer to a node-layer tuple and the term *physical node* to refer to the collection of all state nodes that represent the same entity.

differentiated using heterogeneous spreading rates). The elements of the transition tensor for a classical random walk on a multilayer network are thus

$$P_{j\beta}^{i\alpha} = \frac{A_{j\beta}^{i\alpha}}{\sum_{j\beta \in V_M} A_{j\beta}^{i\alpha}}. \tag{4}$$

An alternative way to generalize the concept of a random walk is by using a *physical random walk* (De Domenico et al., 2014; De Domenico et al., 2015) with transition-tensor elements

$$\tilde{P}_{j\beta}^{i\alpha} = \frac{A_{i\beta}^{i\alpha}}{\sum_{\beta \in L} A_{i\beta}^{i\alpha}} \frac{A_{j\beta}^{i\beta}}{\sum_{j \in V} A_{j\beta}^{i\beta}}. \tag{5}$$

One time step of this physical random walk corresponds to a random walker first switching layers with probabilities proportional to the weights of the interlayer edges and then performing an ordinary random-walk step in the new layer.<sup>2</sup> This type of physical random walk on a multilayer network is equivalent to a classical random walk on a transformed multilayer network with adjacency-tensor elements  $\tilde{A}_{j\beta}^{i\alpha} = A_{i\beta}^{i\alpha} A_{j\beta}^{i\beta}$ . This transformed multilayer network has non-diagonal, directed interlayer edges even when the original multilayer network is undirected.

In this paper, we consider two types of random walks that have been proposed to study communities in multiplex networks: the classical random walk with uniform categorical coupling (Mucha et al., 2010) and the relaxed random walk (De Domenico et al., 2015). For the classical random walk with uniform categorical coupling (which we henceforth call the “classical random walk” for short), we introduce interlayer edges with homogeneous weight  $\omega \in [0, \infty)$  between all pairs of state nodes that correspond to the same physical node. That is, we define  $A_{i\beta}^{i\alpha} = \omega$ ,  $\alpha \neq \beta$  in Equation (4). For the relaxed random walk, we constrain the random walker to follow an edge within the same layer with probability  $1 - r$  (so  $r \in [0, 1]$ ) and allow it to choose uniformly at random among all intralayer edges attached to the same physical node with probability  $r$ . Thus, the transition tensor for the relaxed random walk has elements

$$P_{j\beta}^{i\alpha} = (1 - r)\delta(\alpha, \beta) \frac{A_{j\alpha}^{i\alpha}}{\sum_{j \in V} A_{j\alpha}^{i\alpha}} + r \frac{A_{j\beta}^{i\beta}}{\sum_{j \in V, \beta \in L} A_{j\beta}^{i\beta}}. \tag{6}$$

Alternatively, one can think of the relaxed random walk as a physical random walk with the interlayer edges defined as

$$A_{i\beta}^{i\alpha} = (1 - r)\delta(\alpha, \beta) \sum_{j \in V, \alpha \in L} A_{j\alpha}^{i\alpha} + r \sum_{j \in V} A_{j\beta}^{i\beta}.$$

In Figure 1, we illustrate the classical and relaxed random walks. In Section 2, we discuss how we use random walks to identify communities and characterize mesoscale structures in multilayer networks.

<sup>2</sup> Note that this definition of a physical random walk assumes that the multilayer network has diagonal coupling (i.e., that all interlayer edges are between state nodes that represent the same node).




Fig. 1. Illustration of two types of random walks on a multilayer network. The walks differ in the way that random walkers change layers. (a) *Classical random walk* (Mucha et al., 2010; De Domenico et al., 2014), in which we introduce interlayer edges with weight  $\omega$  between state nodes (i.e., node-layer tuples) that represent the same physical node in adjacent layers. (b) *Relaxed random walk* (De Domenico et al., 2015), in which a random walker is constrained at each step to follow an edge within the same layer with probability  $(1 - r)$  and can choose any intralayer edge attached to the same physical node in any layer (including its own) with probability  $r$ . In the latter case, the walker chooses uniformly at random from the set of all neighbors across all layers.

## 2 Local communities and network community profiles

We seek to contrast the behavior of different types of random walks on multilayer networks. From a dynamical perspective, communities correspond to sets of state nodes that create bottlenecks to a diffusive dynamical process. For a random walk, one measure to quantify bottlenecks is the conductance<sup>3</sup> (Jerrum & Sinclair, 1988)

$$\phi(S) = \frac{\sum_{i\alpha \in S} \sum_{j\beta \notin S} P_{j\beta}^{i\alpha} P_{i\alpha}(\infty)}{\sum_{i\alpha \in S} P_{i\alpha}(\infty)} \tag{7}$$

of a set  $S \subset V_M$  of state nodes. Conductance measures the outflow of random walkers from a set of state nodes relative to the total number of random walkers within the set at stationarity. If a set of state nodes constitutes a bottleneck to a random walk, only a few of the random walkers present within the set should leave the set in a given time step, so the set should have low conductance. The two extreme cases are  $\phi(S) = 1$  if  $S$  has no internal flow (i.e., no state node in  $S$  is adjacent to any other state node in  $S$ ) and  $\phi(S) = 0$  if  $S$  is disconnected from the rest of a network.

Different types of random walks correspond to different notions of what it means for a set of nodes to be a good community. Our choice of conductance as a measure of community quality is motivated by its nice theoretical properties. The presence of low-conductance sets (i.e., sets that are considered to be “good” communities based on the conductance measure) relates directly to slow mixing of a random walk (Mihail, 1989). There are efficient algorithms for identifying low-conductance sets with known approximation guarantees (Andersen et al., 2006; Leskovec et al., 2009). There is also some empirical evidence that conductance is

<sup>3</sup> For a random walk on an undirected, single-layer network, this definition of conductance is equivalent to the conductance in Leskovec et al. (2009) and Jeub et al. (2015).

an effective measure for evaluating community quality and that other measures for evaluating community quality give similar results in practice (Yang & Leskovec, 2015). However, conductance also has some limitations as a measure of community quality. Most notably, it is not very sensitive to the internal connectivity of putative communities. In the most extreme case, low-conductance sets may even be internally disconnected (Leskovec et al., 2009, 2010; Jeub et al., 2015). Our choice of algorithm for identifying local communities (see our discussion below) somewhat mitigates this problem, as it implicitly optimizes the internal connectivity of the identified communities (Leskovec et al., 2010).

We use the ACLCUT method (Andersen et al., 2006; Leskovec et al., 2009; Jeub et al., 2015) to identify putative communities. The ACLCUT method is based on locally ranking nodes near a seed node by approximating a personalized PageRank (PPR) score. Given an appropriate random walk (or other Markov process<sup>4</sup>), one can define the associated PPR score of state node  $ix$  as the solution to the equation

$$\text{PPR}(s, \gamma)_{ix} = \gamma \sum_{j\beta \in V_M} P_{ix}^{j\beta} \text{PPR}(s, \gamma)_{j\beta} + (1 - \gamma)s_{ix}, \quad (8)$$

where  $s$  is a probability distribution that determines the seed nodes for the method (Gleich, 2015) and  $\gamma \in [0, 1]$  is a teleportation parameter. We use two different types of seeding procedure with the ACLCUT method:

- seeding using a state node  $j\beta$ , where  $s_{ix} = \delta(ix, j\beta)$  in Equation (8); and
- seeding using a physical node  $j$ , where  $s_{ix} = \delta(i, j)/|j|$  in Equation (8).

We describe the ACLCUT method in more detail in Appendix A.

Our main tool for summarizing size-resolved community structure is a network community profile (NCP) (Leskovec et al., 2009). An NCP shows the quality of the best community of a given size (i.e., number of nodes) as a function of community size. Because we are using conductance as a measure of community quality, we define the “best” community as the one with the lowest conductance. Hence, we define the NCP as

$$\text{NCP}(k) = \min_{S \subset V_M, |S|=k} \phi(S). \quad (9)$$

We also use local NCPs, where we constrain the communities to contain a given seed set  $S_0$  of state nodes. That is,

$$\text{localNCP}(k, S_0) = \min_{S \subset V_M, |S|=k, S_0 \subset S} \phi(S). \quad (10)$$

An NCP of a network can reveal interesting structural features about the network. In particular, its qualitative shape can reveal the global organization of a network (Jeub et al., 2015). A local NCP is useful for identifying communities at different scales associated with a particular seed node (or seed set of nodes).

Our code for identifying local communities and visualizing networks is available at <https://github.com/LJeub>. There also exists a recently proposed extension of

<sup>4</sup> More generally, it would also be both fruitful and interesting to develop local community-detection methods using dynamical processes that are not Markovian. A good start would be to use our approach through suitable adaptations of other processes that have been used to examine community structure in networks. Examples include Kuramoto phase oscillators (Arenas et al., 2006); epidemic spreading processes (Ghosh et al., 2014); and higher-order Markovian processes, such as those that have been employed in the study of “memory networks” (Rosvall et al., 2014).

the ACLCUT method (Gleich & Kloster, 2016) that allows one to sample local NCPs more efficiently, although we did not use it for our computations in this paper.

### 3 Synthetic benchmark multiplex networks

We now explore the behavior of the two different random walks on synthetic networks with known, planted community structure. The networks that we consider each have  $n = 1,000$  nodes and  $l = 10$  layers (for a total of 10,000 state nodes, as we assume that every node exists on all layers) and  $c = 10$  planted communities.

We generate the planted community structure in the different layers in the following way. We first generate a background community structure  $S_b$  by sampling the community assignment for each node uniformly at random from  $\{1, \dots, c\}$ . We then generate the planted community structure  $S_p$  for the state nodes such that a state node inherits the background community assignment of the corresponding physical node with probability  $1 - \lambda$  and otherwise its community assignment is sampled uniformly at random from  $\{1, \dots, c\}$ .

Given the community assignments for the state nodes, we sample the intralayer edges for the network independently from a stochastic block model, such that an edge between two state nodes in the same layer and with the same community assignment is present with probability  $p_{in}$  and an edge between two state nodes in the same layer but with different community assignment is present with probability  $p_{out}$ . The ratio  $p_{out}/p_{in}$  determines the strength of the community structure of the benchmark, where a small ratio indicates strong community structure. The parameter  $\lambda$  controls the dependency between the layers; the layers have identical community structure for  $\lambda = 0$ , and community structures in different layers are progressively less related to each other with increasing  $\lambda$ .

In Figures 2 and 3, we illustrate the ability of our local community-detection methods to recover planted structure as one varies  $p_{out}$  and  $\lambda$ . We fix  $p_{in} = 0.1$ ,  $n = 1000$ , and  $l = c = 10$ . We use local NCPs to identify communities in the following way. First, we select a state node  $ix$  uniformly at random as a seed node for the local NCP. We then identify the best community for  $ix$  as the community that achieves the minimum conductance

$$S_{\min}(ix) = \operatorname{argmin}_{S \subset V_M, ix \in S} \phi(S)$$

among all communities that include the state node  $ix$ . This construction throws away a lot of information, as local minima of a local NCP can reveal interesting aspects of community structure. However, examining only globally optimal communities for a seed node enables one to easily compare algorithmically obtained community structure with the planted structure. To compare the performance of the different random walks, we use the Jaccard coefficient (Jaccard, 1912)

$$J(S_1, S_2) = \frac{|S_1 \cap S_2|}{|S_1 \cup S_2|},$$

which measures the similarity between two sets  $S_1$  and  $S_2$ . (We obtain the same qualitative results when we compute normalized mutual information.) For a given seed node  $ix$ , we compare the algorithmically identified community  $S_{\min}(ix)$  to both its planted community  $S_p(ix)$  (i.e., the set of all state nodes with the same planted



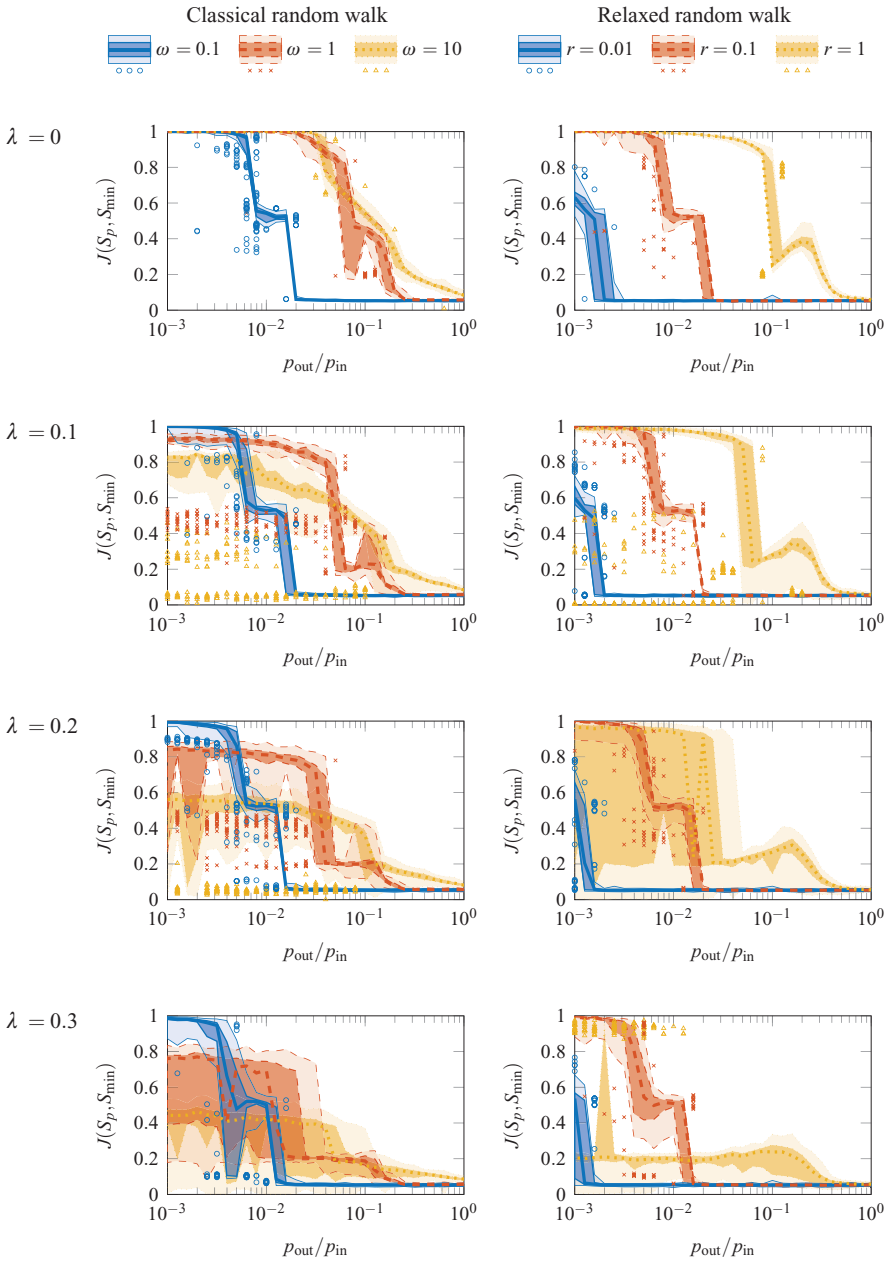


Fig. 2. Recoverability of planted community structure using local NCPs. For each value of  $p_{out}/p_{in}$ , we sample local NCPs for 100 uniformly random seed nodes and compare the best community  $S_{min}$  identified by the local NCP to the planted community  $S_p$  for the seed node using the Jaccard coefficient. The curves indicate the median of the Jaccard-coefficient distributions, the darkly shaded regions indicate the second and third quartiles of the distributions, and the lightly shaded regions indicate the bulks of the distributions. Markers indicate outliers. (Color online)



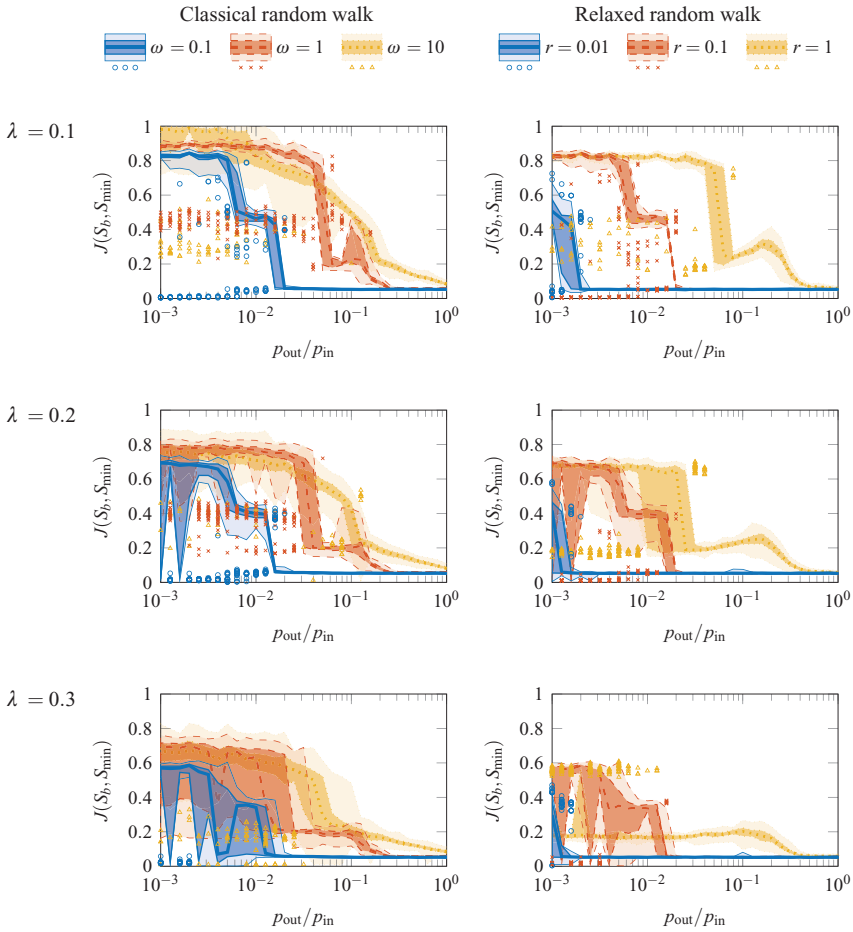


Fig. 3. Recoverability of background community structure using local NCPs. For each value of  $p_{out}/p_{in}$ , we sample local NCPs for 100 uniformly random seed nodes and compare the best community  $S_{min}$  identified by the local NCP to the background community  $S_b$  for the seed node using the Jaccard coefficient. For  $\lambda = 0$  (not shown), planted and background community structure are identical. The curves indicate the median of the Jaccard-coefficient distributions, the darkly shaded regions indicate the second and third quartiles of the distributions, and the lightly shaded regions indicate the bulks of the distributions. Markers indicate outliers. (Color online)

community assignment as  $i\alpha$ ) and its background community  $S_b(i\alpha)$  (i.e., the set of all state nodes whose corresponding physical node has the same background community assignment as the physical node  $i$ ). In Figure 2, we show the distribution of the Jaccard coefficient  $J(S_p(i\alpha), S_{min}(i\alpha))$  between the planted community and best identified community for uniformly random samples of 100 seed nodes. We show the distribution of the Jaccard coefficient  $J(S_b(i\alpha), S_{min}(i\alpha))$  between the background community and best identified community for the same samples in Figure 3.

As one can see from Figure 2, the performances of the two different random walks are comparable to each other; neither one is clearly better than the other. One interesting result is that, as we increase  $\lambda$ , it is increasingly pronounced that there are “good” and “bad” seed nodes for identifying community structure. For  $\lambda = 0$ ,

Table 1. Example network data sets.

	Nodes	Edges	Layers
European Airline Network (Cardillo et al., 2013)	450 airports	3558 (undirected, unweighted)	37 airlines
Lazega Law Firm Network (Lazega & Pattison, 1999; Lazega, 2001; Snijders et al., 2006)	71 employees	2223 (directed, unweighted)	3 (advice, friendship, co-work)

the variability in the Jaccard coefficient for different seed nodes is fairly small, but as we increase  $\lambda$ , the number of outliers increases and we observe increasing variability in the distribution of the Jaccard coefficients.

For small  $\lambda$ , strong interlayer coupling (i.e., large values of  $\omega$  or  $r$ , so that the rate of switching layers is high) helps identify the planted partition. For large enough values of  $\lambda$ , there is a range of  $p_{\text{out}}$  for which community structure is sufficiently strong that random walks with weaker interlayer coupling can outperform those with stronger interlayer coupling. This is already true when  $\lambda = 0.1$ , and the difference becomes more pronounced as one increases  $\lambda$ .

For strong interlayer coupling (in particular, for  $\omega = 10$  and  $r = 1$ ), the two types of random walks lose their ability to identify the planted structure as we increase  $\lambda$  in rather different ways. For small values of  $\lambda$  (in particular, for  $\lambda = 0.1$ ), the classical random walk identifies the background community structure (see Figure 3) rather than the planted community structure. As one increases  $\lambda$  further, its performance at detecting both background and planted community structure decreases gradually for all seed nodes.

In contrast, the performance of the relaxed walk deteriorates in a different way. It loses the ability to identify the planted structure for progressively more choices of seed nodes as we increase  $\lambda$ , but it still performs remarkably well for some seed nodes even when  $\lambda = 0.3$ . For all combinations of  $r$  and  $\lambda$  that we tested, the communities that we obtain from the relaxed walk resemble the planted communities more closely than they resemble the background communities.

#### 4 Empirical multiplex networks

We now illustrate our methodology on two empirical multiplex networks.<sup>5</sup> Our first example is the European Airline Network (Cardillo et al., 2013), a multiplex transportation network with 37 layers, where each layer includes the flights for a single airline. Our second example is the Lazega Law Firm Network (Lazega & Pattison, 1999; Lazega, 2001; Snijders et al., 2006), a multiplex social network with three layers that represent advice, friendship, and co-work relationships between partners and associates of a corporate law firm. In Table 1, we highlight some key properties of these two networks.

<sup>5</sup> Note that for many multilayer networks (and, in particular, for the example networks that we examine in this paper), data on the weights of interlayer edges are not explicitly available (Kivelä et al., 2014).

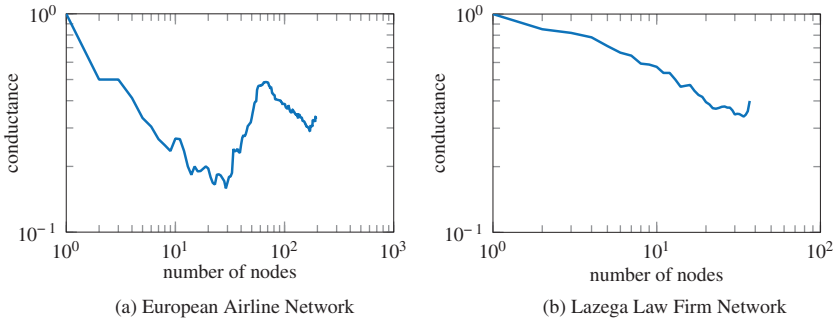


Fig. 4. *Network community profiles (NCPs) of two aggregated empirical networks.* We plot the quality (as measured by conductance) of the best community of each size (as measured by the number of nodes that are a member of the community). (a) The NCP of the aggregated European Airline Network has a shape that one sees in networks with a core–periphery structure. (b) The NCP of the aggregated Lazega Law Firm Network is slightly downward-sloping, and the high minimum conductance indicates that the aggregated network has no clear communities. (Color online)

In Figure 4, we show the NCPs of aggregated networks that we construct from our example multiplex networks. We define the weight of an edge between two nodes in an aggregated network as the number of edges between associated state nodes in the corresponding multilayer network. That is, the adjacency matrix  $\bar{A}$  of the aggregated network has entries

$$\bar{A}_{ij} = \sum_{\alpha \in L} A_{i\alpha}^{j\alpha}.$$

The plots in Figure 4 give a point of comparison for the NCPs of the multilayer networks (see Figures 5 and 6). In its aggregated form, the European Airline Network has an NCP that is suggestive of a core–periphery structure (see Csermely et al., 2013 for a review of such structure), although one cannot conclude this with certainty because the network is small and the conductance values are large. We do not observe any clear structure (and, in particular, no clear community structure) in the Lazega Law Firm Network.

In Figures 5 and 6, we explore the multilayer structure of the airline and law-firm networks. An interesting aspect of multilayer networks is that one can use either physical nodes or state nodes as seeds to sample local communities. We compare the results of these two sampling approaches in Figures 5 and 6. In these two networks, the two approaches produce very similar results. In some cases, sampling using physical nodes can result in slightly better communities. (See the thick and thin solid blue curves in Figure 5(a) for community sizes between  $10^2$  to  $10^3$  state nodes.) In other cases, sampling using physical nodes results in slightly worse communities. (See the thick and thin solid blue curves and dashed red curves in panels (a) and (b) of Figure 6.)

For directed networks (e.g., the network in Figure 6), random walks are not necessarily ergodic, so they might not have a unique stationary distribution. We use so-called “unrecorded edge teleportation”<sup>6</sup> (Lambiotte & Rosvall, 2012; De

<sup>6</sup> For “unrecorded teleportation”, one does not include transitions that result from teleportation when computing flow between nodes.

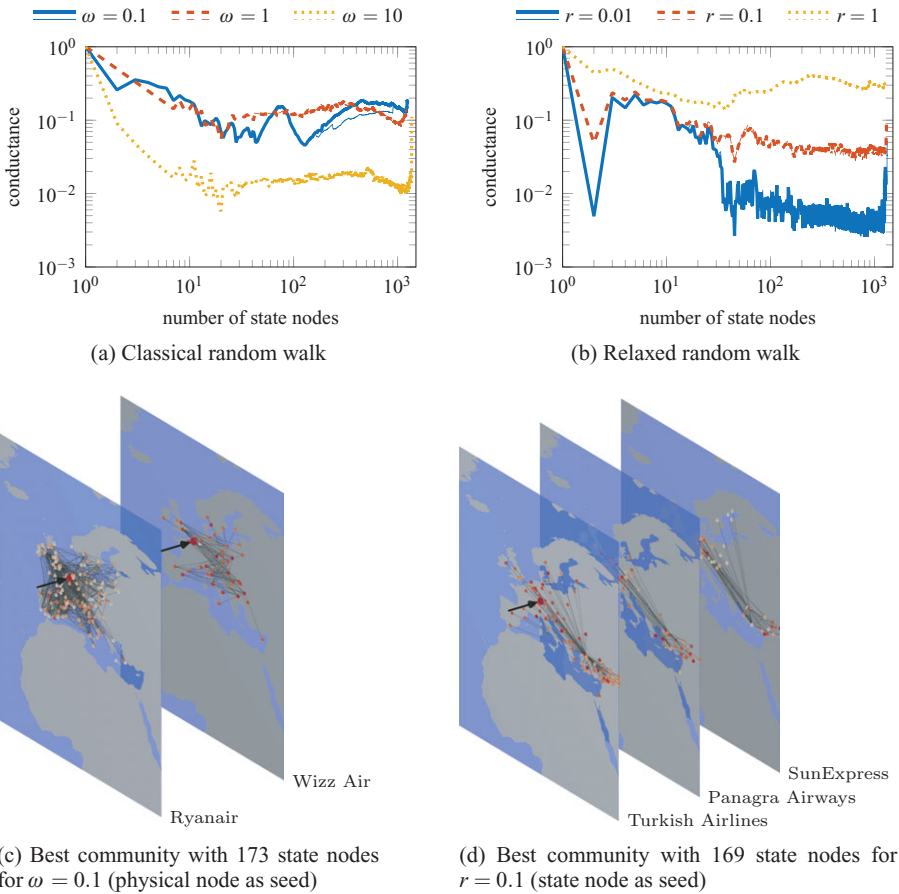


Fig. 5. *European Airline Network*. Panels (a) and (b) show NCPs for this network. We plot the quality (as measured by conductance) of the best community of each size (as measured by the number of state nodes that are a member of the community). Sampling using physical nodes (thin curves) versus using state nodes (thick curves) leads to very similar results, and the thin curves are typically hidden underneath the thick curves in this example. Panels (c) and (d) illustrate some of the communities that we obtain. We shade the state nodes in a community from dark red to light gray based on the rank of their corresponding entry in the degree-normalized PPR-vector that we use to identify the community. (See Appendix A for details.) The large arrows point to the seed nodes. For small layer-jumping probability  $r$  in the relaxed random walk and small interlayer edge weight  $\omega$  in the classical random walk, the best communities tend to consist of sets of similar types of airlines (e.g., they fly to the same airport, are low-cost airlines, or share some other feature). Even for  $r = 1$ , the relaxed walk still predominantly identifies this type of community. By contrast, for large values of  $\omega$ , the classical random walk finds relatively geographically localized communities. The prominent dips in the NCPs in panel (b) for communities consisting of two state nodes are the result of a spurious connection in the data set that creates a bottleneck for the relaxed random walk. (Color online)

Domenico et al., 2015) to ensure that the random walk is ergodic. This corresponds to replacing the stationary distribution of the random walk in the definition of conductance in Equation (7) by a PPR vector in which the seed vector is proportional

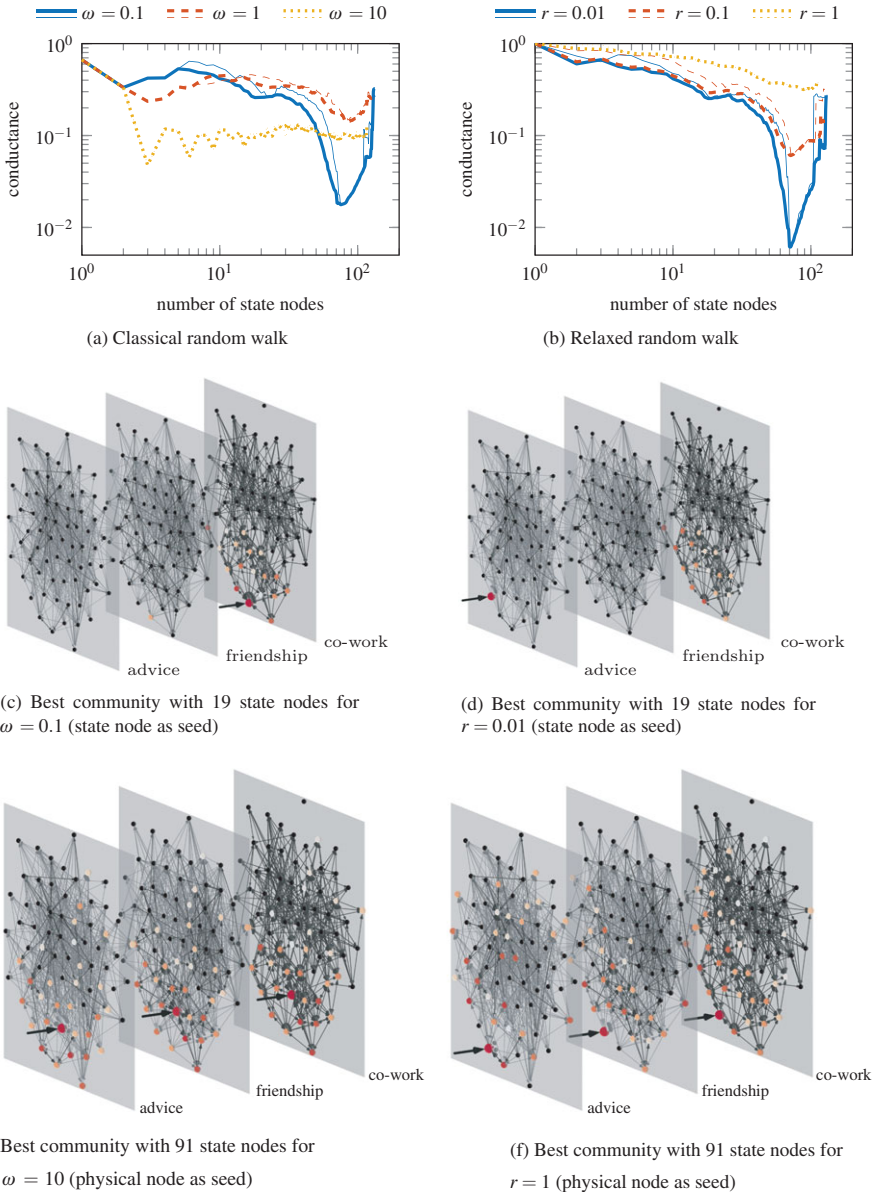


Fig. 6. *Lazega Law Firm Network*. Panels (a) and (b) show NCPs for this network. As with the European Airline Network, when we use a small layer-jumping probability  $r$  in the relaxed random walk and a small interlayer edge weight  $\omega$  in the classical random walk, we obtain similar results even though we consider two different dynamical processes. We also again obtain similar results whether we use a state node or a physical node as a seed. For both types of random walks, with our choice of interlayer connection probability, the communities tend to be localized to a single layer. The prominent minimum in the NCPs at 71 nodes is the result of a community that contains all state nodes in the “co-work” layer. The communities that we highlight in panels (c) and (d) are responsible for the other, less-pronounced minima in the NCPs at 19 nodes. They contain the members of the firm who are based in the Hartford office. For large  $\omega$ , the classical random walk yields communities that are largely “coherent” across layers: If a state node is a member of a community, then the other state nodes associated with the same physical node also tend to be in that community. In contrast, communities from the relaxed random walk are not as coherent across layers. (Color online)

to the vector of in-degrees of the state nodes. That is,

$$p_{ix}(\infty) = \text{PPR}(s, \gamma), \quad s_{ix} = \frac{\sum_{j\beta \in V_M} A_{ix}^{j\beta}}{\sum_{ix \in V_M} \sum_{j\beta \in V_M} A_{ix}^{j\beta}}.$$

For our results in Figure 6, we use  $\gamma = 0.95$  (i.e., a teleportation rate of 0.05). For weighted networks in which intralayer weights are very different in different layers, one may also need to rescale the edge weights appropriately (Cranmer et al., 2015) to obtain results that are not dominated by a single layer (or a small set of layers). However, this issue does not arise in our example networks in the present paper.

As one can see from Figures 5 and 6, the NCPs for the multilayer networks look very different from those of the aggregated networks in Figure 4. The NCPs for the relaxed walk with rate  $r = 1$  are exceptions; their shapes resemble those of the NCPs for the aggregated networks. For each of these networks, the multilayer structure introduces bottlenecks to the spreading of the random walks that are not present in the associated aggregated networks. This is also reflected in the types of communities that underlie these bottlenecks.

For the airline network, the best communities identified by random walks with weak interlayer coupling tend to include all state nodes from a given layer or a set of layers. Effectively, the local communities are identifying sets of similar airlines that share many common destinations. However, as we show in Figure 5, the exact communities identified by the classical random walk and relaxed random walk can be very different. Once the interlayer coupling becomes sufficiently strong, the nature of the communities identified by the classical random walk changes completely. The best communities identified by the classical random walk with strong interlayer coupling tend to be localized geographically and span all layers. The relaxed random walk, however, still predominately identifies sets of airlines even when  $r = 1$  (i.e., when the interlayer coupling is maximal). In fact, the communities identified by the relaxed random walk with  $r = 1$  are often rather similar to those identified with  $r = 0.1$ .

For the Lazega Law Firm Network, the layer structure also results in bottlenecks to the random walks when interlayer coupling is weak. This is the cause of the sharp minima in the NCPs in Figure 6 for community sizes of 71 state nodes. At smaller community sizes, one can identify intralayer structures—most notably, a community in the “co-work” layer that consists of the members of the firm that work in the Hartford office. Unlike for the airline network, in the case of the Lazega Law Firm Network, both types of random walk predominantly identify communities that span all layers when the interlayer coupling is sufficiently strong. However, the two types of random walks explore the law-firm network in rather different ways. The classical random walk explores the different layers of the network in a “coherent” manner when the interlayer coupling is strong. That is, if a state node associated with a particular physical node is included in a community, then the other state nodes of that physical node (i.e., its manifestation in the other layers) tend to also belong to the community. However, as we illustrate in panels (e) and (f) of Figure 6, the communities identified by the relaxed random walk tend to be less coherent across layers than those identified by the classical random walk.



To understand the difference in behavior between the relaxed and classical random walks at high layer-switching rates (i.e., for  $r \approx 1$  and for large  $\omega$ ), it is important to examine the behavior of the associated dynamical processes that are induced on the aggregated networks. The dynamical process induced on the aggregated network by the relaxed random walk with  $r = 1$  is simply a standard random walk. However, the process induced by the classical random walk for large  $\omega$  explores the aggregated network much more slowly than a standard random walk, as most of the transitions are between state nodes that represent the same physical node. This results in a downward shift of the NCPs as one increases  $\omega$ . This has a similar effect as introducing a self-loop at each node. (As discussed in Arenas et al. (2008), introducing self-loops is one way to introduce a resolution parameter in the modularity quality function.)

## 5 Discussion and conclusions

We have seen using example synthetic networks that bottlenecks of random walks on a multilayer network can reveal non-trivial multiplex community structure in which community structure in different layers of a network differs but is not independent. We explored two types of random walks—a classical random walk and a relaxed random walk—for identifying structure in our synthetic benchmarks, and we found that they can behave rather differently from each other in some situations. Consequently, different random walks give different community structures in a network, and one thus also expects to observe (although we did not test this directly) different community structures in different global methods (e.g., based on optimizing a quality function) based on the two different types of random walks. Similar results have been noted previously in other contexts (Lambiotte et al., 2009, 2015).

As we saw in Section 4, the behavior of the random walks on a multilayer network is in general very different from the behavior of a random walk on a corresponding aggregated network. Consequently, examining a multilayer network can reveal important information that is not visible in a corresponding aggregated network. In particular, bottlenecks to random walks on a multilayer network can reveal structures, such as sets of related layers and communities confined to a single layer, that are impossible to see in aggregated networks. Although we only considered such a scenario in the context of synthetic networks in the present paper, it is also possible for a multilayer network to possess similar communities across layers. They can span many (or even all) layers (Peixoto, 2015), and then one expects to observe significant bottlenecks for multilayer dynamics with strong interlayer coupling and for dynamics on the aggregated networks (which are monoplex networks, as in Leskovec et al., 2009 and Jeub et al., 2015).

Our approach is very general, and a suite of other dynamical processes can also be used to develop a diverse family of local community-detection methods. In addition to considering other processes, in advancing our work further, it will also be interesting to exploit transformations between ordinary random walks and other types of random walks (Lambiotte et al., 2011; Yan et al., 2016). Another interesting extension of our approach would be to use it as the seed-set-expansion part of a seed-centric algorithm (Kanawati, 2014; Hmimida & Kanawati, 2015; Whang et al., 2016) for detecting communities in multilayer networks.



### Acknowledgments

LGSJ acknowledges a CASE studentship award from the EPSRC (BK/10/039), and LGSJ and MAP were supported in part from the James S. McDonnell Foundation 21st Century Science Initiative Complex Systems Scholar Award grant # 220020177 and the FET-Proactive project PLEXMATH (FP7-ICT-2011-8; grant # 317614) funded by the European Commission. MAP was also supported by the EPSRC (EP/J001759/1). MWM acknowledges funding from the Army Research Office and from the Defense Advanced Research Projects Agency. PJM was supported from the James S. McDonnell Foundation 21st Century Science Initiative Complex Systems Scholar Award grant # 220020315. MAP also thanks SAMSI for supporting several visits and MWM for his hospitality during a sabbatical at Stanford.

### References

- Andersen, R., Chung, F. R. K., & Lang, K. J. (2006). Local graph partitioning using PageRank vectors. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science*. New York, NY, USA: IEEE, pp. 475–486.
- Arenas, A., Díaz-Guilera, A., & Pérez-Vicente, C. J. (2006). Synchronization reveals topological scales in complex networks. *Physical Review Letters*, **96**(11), 114102.
- Arenas, A., Fernández, A., & Gómez, S. (2008). Analysis of the structure of complex networks at different resolution levels. *New Journal of Physics*, **10**(5), 053039.
- Boccaletti, S., Bianconi, G., Criado, R., del Genio, C. I., Gómez-Gardenes, J., Romance, M., . . . Zanin, M. (2014). The structure and dynamics of multilayer networks. *Physics Reports*, **544**(1), 1–122.
- Cardillo, A., Gómez-Gardeñes, J., Zanin, M., Romance, M., Papo, D., del Pozo, F., & Boccaletti, S. (2013). Emergence of network features from multiplexity. *Scientific Reports*, **3**, 1344.
- Coscia, M., Giannotti, F., & Pedreschi, D. (2011). A classification for community discovery methods in complex networks. *Statistical Analysis and Data Mining*, **4**(5), 512–546.
- Cranmer, S. J., Menninga, E. J., & Mucha, P. J. (2015). Kantian fractionalization predicts the conflict propensity of the international system. *Proceedings of the National Academy of Sciences of the United States of America*, **112**(38), 11812–11816.
- Csermely, P., London, A., Wu, L.-Y., & Uzzi, B. (2013). Structure and dynamics of core-periphery networks. *Journal of Complex Networks*, **1**(2), 93–123.
- De Domenico, M., Solè-Ribalta, A., Gómez, S., & Arenas, A. (2014). Navigability of interconnected networks under random failures. *Proceedings of the National Academy of Sciences of the United States of America*, **111**(23), 8351–8356.
- De Domenico, M., Lancichinetti, A., Arenas, A., & Rosvall, M. (2015). Identifying modular flows on multilayer networks reveals highly overlapping organization in social systems. *Physical Review X*, **5**(1), 011027.
- De Domenico, M., Granell, C., Porter, M. A. & Arenas, A. (2016). The physics of spreading processes in multilayer networks. *Nature Physics*, **12**(10), 901–906.
- Delvenne, J.-C., Yaliraki, S. N., & Barahona, M. (2010). Stability of graph communities across time scales. *Proceedings of the National Academy of Sciences of the United States of America*, **107**(29), 12755–12760.
- Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, **486**(3–5), 75–174.
- Fortunato, S., & Hric, D. (2016). Community detection in networks: A user guide. *Physics Reports*, **659**, 1–44.
- Ghosh, R., Lerman, K., Teng, S.-H., & Yan, X. (2014). The interplay between dynamics and networks: Centrality, communities, and Cheeger inequality. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '14. New York, NY, USA: ACM, pp. 1406–1415.
- Gleich, D. F. (2015). PageRank beyond the Web. *SIAM Review*, **57**(3), 321–363.

- Gleich, D. F., & Kloster, K. (2016). Seeded PageRank solution paths. *European Journal of Applied Mathematics*, **27**(6), 812–845.
- Hmimida, M., & Kanawati, R. (2015). Community detection in multiplex networks: A seed-centric approach. *Networks and Heterogeneous Media*, **10**(1), 71–85.
- Holme, P. (2015). Modern temporal network theory: A colloquium. *The European Physical Journal B*, **88**(9), 234.
- Holme, P., & Saramäki, J. (2012). Temporal networks. *Physics Reports*, **519**(3), 97–125.
- Jaccard, P. (1912). The distribution of the flora in the alpine zone. *New Phytologist*, **11**(2), 37–50.
- Jerrum, M., & Sinclair, A. (1988). Conductance and the rapid mixing property for Markov chains: The approximation of the permanent resolved. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*. New York, NY, USA: ACM, pp. 235–244.
- Jeub, L. G. S., Balachandran, P., Porter, M. A., Mucha, P. J., & Mahoney, M. W. (2015). Think locally, act locally: Detection of small, medium-sized, and large communities in large networks. *Physical Review E*, **91**(1), 012821.
- Kanawati, R. (2014). Seed-centric approaches for community detection in complex networks. In G. Meiselwitz (Ed.), *Proceedings of the 6th International Conference on Social Computing and Social Media, SCSM 2014*. Lecture Notes in Computer Science, vol. 8531. Cham, Switzerland: Springer International Publishing, pp. 197–208.
- Kivelä, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., & Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks*, **2**(3), 203–271.
- Kloumann, I., Ugander, J., & Kleinberg, J. (2016). Block models and personalized PageRank. arXiv:1607.03483.
- Kuncheva, Z., & Montana, G. (2015). Community detection in multiplex networks using locally adaptive random walks. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. ASONAM '15. New York, NY, USA: ACM, pp. 1308–1315.
- Lambiotte, R., & Rosvall, M. (2012). Ranking and clustering of nodes in networks with smart teleportation. *Physical Review E*, **85**(5), 056107.
- Lambiotte, R., Delvenne, J.-C., & Barahona, M. (2009). *Laplacian dynamics and multiscale modular structure in networks*. arXiv:0812.1770v3.
- Lambiotte, R., Sinatra, R., Delvenne, J.-C., Evans, T. S., Barahona, M., & Latora, V. (2011). Flow graphs: Interweaving dynamics and structure. *Physical Review E*, **84**(1), 017102.
- Lambiotte, R., Delvenne, J.-C., & Barahona, M. (2015). Random walks, Markov processes and the multiscale modular organization of complex networks. *Transactions on Network Science and Engineering*, **1**(2), 76–90.
- Lazega, E. (2001). *The Collegial Phenomenon: The Social Mechanisms of Cooperation Among Peers in a Corporate Law Partnership*. Oxford, UK: Oxford University Press.
- Lazega, E., & Pattison, P. E. (1999). Multiplexity, generalized exchange and cooperation in organizations: A case study. *Social Networks*, **21**(1), 67–90.
- Leskovec, J., Lang, K. J., Dasgupta, A., & Mahoney, M. W. (2009). Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, **6**(1), 29–123.
- Leskovec, J., Lang, K. J., & Mahoney, M. W. (2010). Empirical comparison of algorithms for network community detection. In *Proceedings of the 19th International Conference on World Wide Web*. New York, NY, USA: ACM, pp. 1308–1315.
- Mihail, M. (1989). Conductance and convergence of Markov chains — A combinatorial treatment of expanders. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*. New York, NY, USA: IEEE, pp. 526–531.
- Mucha, P. J., Richardson, T., Macon, K., Porter, M. A., & Onnela, J.-P. (2010). Community structure in time-dependent, multiscale, and multiplex networks. *Science*, **328**(5980), 876–878.
- Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford, UK: Oxford University Press.
- Peixoto, T. P. (2015). Inferring the mesoscale structure of layered, edge-valued, and time-varying networks. *Physical Review E*, **92**(4), 042807.

- Porter, M. A., Onnela, J.-P., & Mucha, P. J. (2009). Communities in networks. *Notices of the American Mathematical Society*, **56**(9), 1082–1097, 1164–1166.
- Rosvall, M., & Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences of the United States of America*, **105**(4), 1118–1123.
- Rosvall, M., Esquivel, A. V., Lancichinetti, A., West, J. D., & Lambiotte, R. (2014). Memory in network flows and its effects on spreading dynamics and community detection. *Nature Communications*, **5**, 4630.
- Salehi, M., Sharma, R., Marzolla, M., Magnani, M., Siyari, P., & Montesi, D. (2015). Spreading processes in multilayer networks. *IEEE Transactions on Network Science and Engineering*, **2**(2), 65–83.
- Snijders, T. A. B., Pattison, P. E., Robins, G. L., & Handcock, M. S. (2006). New specifications for exponential random graph models. *Sociological Methodology*, **36**(1), 99–153.
- Wasserman, S., & Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge, UK: Cambridge University Press.
- Whang, J. J., Gleich, D. F., & Dhillon, I. S. (2016). Overlapping community detection using neighborhood-inflated seed expansion. *IEEE Transactions on Knowledge and Data Engineering*, **28**(5), 1272–1284.
- Yan, X., Teng, S.-H., Lerman, K., & Ghosh, R. (2016). Capturing the interplay of dynamics and networks through parameterizations of Laplacian operators. *PeerJ Computer Science*, **2**, e57.
- Yang, J., & Leskovec, J. (2015). Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, **42**(1), 181–213.

### Appendix A. Sampling network community profiles (the ACLcut method)

We use the ACLCUT method (Andersen et al., 2006; Leskovec et al., 2009; Jeub et al., 2015) to sample local communities and NCPs. In this appendix, we briefly discuss how we apply this procedure to identify communities using a random walk on a multilayer network.

The key idea behind the ACLCUT method is the use of a “push” procedure (Andersen et al., 2006), which pushes probability mass from the residual vector  $e$  to the PageRank vector  $p$  while preserving the invariant  $p = \text{PPR}(s - e, \gamma)$ . We describe the different parts of the ACLCUT method in Figures A1–A3. In addition

---

```

def ACLCUT( $P, v, S_0, \gamma, \epsilon$ ):
    for  $i\alpha \in V_M$ :                               set up seed vector
        if  $i\alpha \in S_0$ :
             $s_{i\alpha} = 1/|S_0|$ 
        else:
             $s_{i\alpha} = 0$ 
     $p = \text{APPR}(P, v, s, \gamma, \epsilon)$              compute  $\epsilon$ -approximate PageRank vector
     $\phi, \mathcal{S} = \text{SWEEP CUT}(P, v, p/v)$          normalized sweep cut
    return  $\phi, \mathcal{S}$                                return conductance and communities

```

---

Fig. A1. ACLCUT method for sampling local communities. The inputs are the transition tensor  $P$  of the random walk, a seed set  $S_0$  of state nodes, and a vector  $v$  of node volumes. The vector is proportional either to the stationary distribution of  $P$  or (when considering teleportation) to a PageRank vector. The resolution of the method is controlled by the teleportation parameter  $\gamma$  and the truncation parameter  $\epsilon$ .

---

```

def APPR( $P, v, s, \gamma, \epsilon$ ):
     $\tilde{\gamma} = (1 - \gamma)/(1 + \gamma)$                                 convert to equivalent lazy-walk teleportation
    for  $i\alpha \in V_M$ :
         $p_{i\alpha} = 0$                                         initialize PageRank vector
         $e_{i\alpha} = s_{i\alpha}$                                 initialize residual
     $Q = \{i\alpha : e_{i\alpha} \geq \epsilon v_{i\alpha}\}$                 keep track of nodes to update
    while  $Q \neq \emptyset$ :
         $i\alpha = \text{pop}(Q)$                                 select a node to update
         $\tilde{e} = e_{i\alpha}$ 
         $p_{i\alpha} = p_{i\alpha} + \tilde{\gamma}\tilde{e}$                         push probability mass to PageRank vector
         $e_{i\alpha} = (1 - \tilde{\gamma})\tilde{e}/2$ 
        for  $j\beta \in \{j\beta : P_{j\beta}^{i\alpha} > 0\}$ :
             $e_{j\beta} = e_{j\beta} + (1 - \tilde{\gamma})P_{j\beta}^{i\alpha}\tilde{e}/2$ 
         $Q = \{i\alpha : e_{i\alpha} \geq \epsilon v_{i\alpha}\}$                 check  $i\alpha$  and its neighbors to update  $Q$ 
    return  $p$                                             return  $\epsilon$ -approximate PageRank vector

```

---

Fig. A2. Approximate PPR (APPR) procedure for computing  $\epsilon$ -approximate PageRank vectors using only local information. See Fig. A1 for a description of the input arguments.

---

```

def SWEEPCUT( $P, v, p$ ):
     $N = \text{RANKORDER}(V_M, p)$                                 return state nodes in descending order of  $p$ 
     $S = \emptyset$ 
     $c = 0$ 
     $vol = 0$ 
    for  $k \in \{1, \dots, |N|\}$ :
         $i\alpha = N(k)$                                         get the next state node to consider
         $c = c - \sum_{j\beta \in S} P_{i\alpha}^{j\beta} v_{j\beta} + \sum_{j\beta \notin S} P_{j\beta}^{i\alpha} v_{i\alpha}$  update conductance
         $vol = vol + v_{i\alpha}$ 
         $S = S \cup i\alpha$ 
         $\phi_k = c/vol$ 
         $\mathcal{S}(k) = S$ 
    return  $\phi, \mathcal{S}$                                         return conductance values and sweep sets

```

---

Fig. A3. SWEEPCUT procedure for identifying communities based on a ranking vector  $p$  for the state nodes.

to the teleportation parameter  $\gamma$ , the ACLCUT method also depends on a truncation parameter  $\epsilon$ . The ACLCUT method terminates once the residual is small enough so that  $e_{i\alpha} < \epsilon v_{i\alpha}$  for all state nodes  $i\alpha \in V_M$ , where the quantity  $v$  denotes a vector of node volumes. We set  $v = |V_M| \times p(\infty)$ , where  $p(\infty)$  is either the stationary distribution of the random walk or (when considering teleportation) it is a PageRank vector. The rescaling is purely for computational convenience, as from a theoretical perspective the results are invariant under rescaling of the node volumes (because one also rescales  $\epsilon$ ).

To sample an NCP, we use the ACLCUT method to sample communities for different values of  $\epsilon$  and different seed nodes, and we take the lower envelope of the conductance values. To sample a local NCP, we only vary  $\epsilon$  and use the seed set of the local NCP as a seed set for the ACLCUT method. We use 20 logarithmically spaced values for  $\epsilon$  in the interval  $[1/\|v\|_\infty, 1/\|v\|_1]$ , and we fix  $\gamma = 0.998$ . For each

value of  $\epsilon$ , we initially set the sample set  $S$  of potential seed nodes to be either the set of all state nodes (i.e.,  $S = V_M$ ) or the set of all physical nodes (i.e.,  $S = V$ ). We then sample seed nodes uniformly at random without replacement from  $S$  until  $S$  is empty. To avoid excessive computations for small values of  $\epsilon$ , we remove nodes from  $S$  once they have been included in the best local community returned by the ACLCUT method 10 times. This sampling procedure allows one to estimate an NCP in a time that scales almost linearly with the number of state nodes while ensuring good coverage of the structure of a network.