

Reconstruction of Random Graphs With Few Participants

Animashree Anandkumar*, Avinatan Hassidim[†] and Jonathan Kelner[‡]

November 11, 2010

Abstract

We consider the task of topology discovery of sparse random graphs using end-to-end random measurements (e.g., delay) between a subset of nodes, referred to as the participants. The rest of the nodes are hidden, and do not provide any information for topology discovery. We consider topology discovery under two routing models: (a) the participants exchange messages along the shortest paths and obtain end-to-end measurements, and (b) additionally, the participants exchange messages along the second shortest path. For scenario (a), our proposed algorithm results in a sub-linear edit-distance guarantee using a sub-linear number of uniformly selected participants. For scenario (b), we obtain a much stronger result, and show that we can achieve consistent reconstruction when a sub-linear number of uniformly selected nodes participate. This implies that accurate discovery of sparse random graphs is tractable using an extremely small number of participants. Our algorithms are simple to implement, computationally efficient, and exploit the locally tree-like property of sparse random graphs. We finally obtain a lower bound on the number of participants required by any algorithm to reconstruct the original random graph up to a given edit distance. We also demonstrate that while consistent discovery is tractable for sparse random graphs using a small number of participants, in general, there are graphs which cannot be discovered by any algorithm even with a significant number of participants, and with the availability of end-to-end information along all the paths between the participants.

1 Introduction

Inference of global characteristics of large networks using limited local information is an important and a challenging task. The discovery of the underlying network topology is one of the main goals of network inference, and its knowledge is crucial for many applications. For instance, in communication networks, many network monitoring applications rely on the knowledge of the routing topology, e.g., to evaluate the resilience of the network to failures [34, 39]; for network traffic prediction [24, 51] and monitoring [8], anomaly detection [6], or to infer the sources of viruses and rumors in the network [47]. In the context of social networks, the knowledge of topology is useful for inferring many characteristics such as identification of hierarchy and community structure [27],

*A. Anandkumar is with the Center for Pervasive Communications and Computing, Electrical Engineering and Computer Science Dept., University of California, Irvine, USA 92697. Email: a.anandkumar@uci.edu

[†]A. Hassidim is with Google Labs, Israel. Email: avinatanh@gmail.com

[‡]J. Kelner is with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge MA 02139. Email: kelner@mit.edu

prediction of information flow [5, 53], or to evaluate the possibility of information leakage from anonymized social networks [9].

Traditionally, inference of routing topology in communication networks has relied on tools such as traceroute and mtrace [3] to generate path information between a subset of nodes. However, these tools require cooperation of intermediate nodes or routers to generate messages using the Internal Control Message Protocol (ICMP). Increasingly, today many routers block traceroute requests due to privacy and security concerns [30, 54], thereby making inference of topology using traceroute inaccurate. Moreover, traceroute requests are not scalable for large networks, and cannot discover layer-2 switches and MPLS (Multi-protocol Label Switching) paths, which are increasingly being deployed [42].

The alternative approach for topology discovery is the approach of *network tomography*. Here, topology inference is carried out from end-to-end packet probing measurements (e.g., delay) between a subset of nodes, without the need for cooperation between the intermediate (i.e., non-participating) nodes in the network. Due to its flexibility, such approaches are gaining increasing popularity (see Section 1.2 for details).

The approach of topology discovery using end-to-end measurements is also applicable in the context of social networks. In many social networks, some nodes may be unwilling to participate or cooperate with other nodes for discovering the network topology, and there may be many hidden nodes in “hard to reach” places of the network, e.g., populations of drug users, and so on. Moreover, in many networks, there may be a cost to probing nodes for information, e.g., when there is a cash reward offered for filling out surveys. For such networks, it is desirable to design algorithms which can discover the overall network topology using small fraction of participants who are willing to provide information for topology discovery.

There are many challenges to topology discovery. The algorithms need to be computationally efficient and provide accurate reconstruction using a small fraction of participating nodes. Moreover, inference of large topologies is a task of *high-dimensional learning* [52]. In such scenarios, typically, only a small number of end-to-end measurements are available relative to the size of the network to be inferred. It is desirable to have algorithms with low *sample complexity* (see Definition 3), where the number of measurements required to achieve a certain level of accuracy scales favorably with the network size.

It is indeed not tractable to achieve all the above objectives for discovery of general network topologies using an arbitrary set of participants. There are fundamental identifiability issues, and in general, no algorithm will be able to discover the underlying topology. We demonstrate this phenomenon in Section 8.2, where we construct a small network with a significant fraction of participants which suffers from non-identifiability. Instead, it is desirable to design topology discovery algorithms which have guaranteed performance for certain classes of graphs.

We consider the class of Erdős-Rényi random graphs [13]. These are perhaps the simplest as well as the most well-studied class of random graphs. Such random graphs can provide a reasonable explanation for peer-to-peer networks [33] and social networks [41]. We address the following issues in this paper: can we discover random graphs using a small fraction of participating nodes, selected uniformly at random? can we design efficient algorithms with low sample complexity and with provable performance guarantees? what kinds of end-to-end measurements between the participants are useful for topology discovery? finally, given a set of participants, is there a lower bound on the error (edit distance) of topology discovery that is achievable by any algorithm? Our work addresses these questions and also provides insights into many complex issues involved in

topology discovery.

1.1 Summary of Contributions

We consider the problem of topology discovery of sparse random graphs using a uniformly selected set of participants. Our contributions in this paper are three fold. First, we design an algorithm with provable performance guarantees, when only minimal end-to-end information between the participants is available. Second, we consider the scenario with additional information, and design a discovery algorithm with much better reconstruction guarantees. Third, we provide a lower bound on the edit distance of the reconstructed graph by any algorithm, for a given number of participants. Our analysis shows that random graphs can be discovered accurately and efficiently using an extremely small number of participants.

We consider reconstruction of the giant component of the sparse random graph up to its minimal representation, where there are no redundant hidden nodes (see Section 3.1). Our end-to-end measurement model consists of random samples (e.g., delay) along the shortest paths between the participants. Using these samples, we design the first random-discovery algorithm, referred to as the RGD1 algorithm, which performs local tests over small groups of participating nodes (known as the *quartet tests*), and iteratively merges them with the previously constructed structure. Such tests are known to be accurate for tree topologies [12], but have not been previously analyzed for random-graph topologies. We provide a sub-linear edit-distance guarantee under RGD1 using a sub-linear number of participants, with respect to the size of the network. The algorithm is also simple to implement, and is computationally efficient.

We then extend the algorithm to the scenario where additionally, there are end-to-end measurements available along the second shortest paths between the participating nodes. Such information is available since nodes typically maintain information about alternative routing paths, should the shortest path fail. In this scenario, our algorithm RGD2, has a drastic improvement in accuracy under the same set of participating nodes. Specifically, we demonstrate that consistent discovery (i.e., zero edit distance) can be achieved under RGD2 algorithm when there are more than $n^{7/8}$ number of participants, where n is the network size for homogeneous delays (we have a general bound for heterogeneous delays). Thus, we can achieve accurate topology discovery of random graphs using an extremely small number of participants. For both our algorithms, the sample complexity is poly-logarithmic in the network size, meaning that the number of end-to-end measurement samples needs to scale poly-logarithmically in the network size to obtain the stated edit-distance guarantees.

Our analysis in this paper thus reveals that sparse random graphs can be efficiently discovered using a small number of participants. Our algorithms exploit the *locally tree-like* property of random graphs [13], meaning that these graphs contain a small number of short cycles. This enables us to provide performance guarantees for quartet tests which are known to be accurate for tree topologies, and this is done by carefully controlling the distances used by the quartet tests. At the same time, we exploit the presence of cycles in random graphs to obtain much better guarantees than in the case of tree topologies. In other words, while tree topologies require participation of at least half the number of nodes (i.e., the leaves) for accurate discovery, random-graph topologies can be accurately discovered using a sub-linear number of participants.

Finally, we provide lower bounds on the reconstruction error under any algorithm for a given number of participants. Specifically, we show that if less than roughly \sqrt{n} nodes participate in topology discovery, reconstruction is impossible under any algorithm, where n is the network size. We also discuss topology discovery in general networks, and demonstrate identifiability issues in-

volved in the discovery process. We construct a small network with a significant fraction of nodes as participants which cannot be reconstructed using end-to-end information on all possible paths between the participants. This is in contrast to random graphs, where consistent and efficient topology discovery is possible using a small number of participants.

To the best of our knowledge, this is the first work to undertake a systematic study of random-graph discovery using end-to-end measurements between a subset of nodes. Although we limit ourselves to the study of random graphs, our algorithms are based on the locally tree-like property, and are thus equally applicable for discovering other locally tree-like graphs such as the d -regular graphs and the *scale-free* graphs; the latter class is known to be a good model for social networks [41] and peer-to-peer networks [33].

1.2 Related Work

Network tomography has been extensively studied in the past and various heuristics and algorithms have been proposed along with experimental results on real data. For instance, the area of mapping the internet topology is very rich and extensive, e.g., see [1, 2, 4, 24, 29, 31, 48, 50]. In the context of social networks, the work in [37] considers prediction of positive and negative links, the work in [28] considers inferring networks of diffusion and influence and the work in [40] considers inferring latent social networks through spread of contagions. A wide range of network tomography solutions have been proposed for general networks. See [16] for a survey.

Topology discovery is an important component of network tomography. There have been several theoretical developments on this topic. The work in [18] provides hardness results for topology discovery under various settings. Topology discovery under availability of different kinds of queries have been previously considered, such as:

(i) *Shortest path query*, where a query to a node returns all the shortest paths (i.e., list of nodes in the path) from that node to all other nodes [11]. This is the strongest of all queries. These queries can be implemented by using Traceroute on Internet. In [11], the combinatorial optimization problem of selecting the smallest subset of nodes for such queries to estimate the network topology is formulated. The work in [25] considers discovery of random graphs using such queries.

(ii) *Distance query*, where a query to a node returns all the shortest-path distances (instead of the complete list of nodes) from that node to any other node in the network [25]. These queries are available for instance, in Peer-to-Peer networks through the Ping/Pong protocol. This problem is related to the landmark placement, and the optimization problem of having smallest number of landmarks is known as the metric dimension of the graph [35]. The work in [45] considers reconstruction of tree topologies using shortest-path queries.

(iii) *Edge-based queries*: There are several types of edge queries such as detection query, which answer whether there is an edge between two selected nodes, or counting query, which returns number of edges in a selected subgraph [38, 44], or a cross-additive query, which returns the number of edges crossing between two disjoint sets of vertices [17].

However, all the above queries assume that all the nodes (with labels) are known a priori, and that there are no hidden (unlabeled) nodes in the network. Moreover, most of the above works consider unweighted graphs, which are not suitable when end-to-end delay (or other weighted) information is available for topology discovery. As previously discussed, the above queries assume extensive information is available from the queried objects, and this may not be feasible in many networks.

Topology discovery using end-to-end delays between a subset of nodes (henceforth, referred to as participating nodes), has been previously studied for tree topologies using unicast traffic in [12, 42, 49] and multicast traffic [21]. The algorithms are inspired by *phylogenetic* tree algorithms. See [22] for a thorough review. Most of these algorithms are based on a series of local tests known as the *quartet-based distance tests*. Our algorithms are inspired by, and are based on quartet methods. However, these algorithms were previously applied only to tree topologies, and here, we show how algorithms based on similar ideas can provide accurate reconstruction for a much broader class of locally-tree like graphs such as the sparse random graphs.

2 System Model

Notation

For any two functions $f(n), g(n)$, $f(n) = O(g(n))$ if there exists a constant M such that $f(n) \leq Mg(n)$ for all $n \geq n_0$ for some fixed $n_0 \in \mathbb{N}$. Similarly, $f(n) = \Omega(g(n))$ if there exists a constant M' such that $f(n) \geq M'g(n)$ for all $n \geq n_0$ for some fixed $n_0 \in \mathbb{N}$, and $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$. Also, $f(n) = o(g(n))$ when $f(n)/g(n) \rightarrow 0$ and $f(n) = \omega(g(n))$ when $f(n)/g(n) \rightarrow \infty$ as $n \rightarrow \infty$. We use notation $\tilde{O}(g(n)) = O(g(n)\text{poly log } n)$. Let $\mathbb{I}[A]$ denote indicator of an event A .

Let G_n denote a random graph with probability measure \mathbb{P} . Let \mathcal{Q} be a graph property (such as being connected). We say that the property \mathcal{Q} for a sequence of random graphs $\{G_n\}_{n \in \mathbb{N}}$ holds asymptotically almost surely (a.a.s.) if,

$$\lim_{n \rightarrow \infty} \mathbb{P}(G_n \text{ satisfies } \mathcal{Q}) = 1.$$

Equivalently, the property \mathcal{Q} holds for almost every (a.e.) graph G_n .

For a graph G , let $\mathcal{C}(l; G)$ denote the set of (generalized) cycles¹ of length less than l in graph G . For a vertex v , let $\deg(v)$ denote its degree and for an edge e , let $\deg(e)$ denote the total number of edges connected to either of its endpoints (but not counting the edge e). Let $B_R(v)$ denote the set of nodes within hop distance R from a node v and $\Gamma_R(v)$ is the set of nodes exactly at hop distance R . The definition is extended to an edge, by considering union of sets of the endpoints of edge. Denote the shortest path (with least number of hops) between two nodes i, j as $\text{path}(i, j; G)$ and the second shortest path as $\text{path}_2(i, j; G)$. Denote the number of H -subgraphs in G as $N_{H;G}$.

2.1 Random Graphs

We assume that the unknown network topology is drawn from the ensemble of Erdős-Rényi random graphs [13]. This random graph model is arguably the simplest as well the most well-studied model. Denote the random graph as $G_n \in \mathcal{G}(n, c/n)$, for $c < \infty$, where n is the number of nodes and each edge occurs uniformly with probability c/n . This implies a constant average degree of c for each node, and this regime is also known as the “sparse” regime of random graphs.

It is well known that sparse random graphs exhibit a phase transition with respect to the number of components. When $c > 1$, there is a giant component containing $\Theta(n)$ nodes, while all the other components have size $\Theta(\log n)$ [7, Ch. 11]. This regime is known as the *super-critical* regime. On

¹A generalized cycle of length l is a connected graph of l nodes with l edges. A cycle refers to a generalized cycle unless otherwise mentioned.

the other hand, when $c < 1$, there is no giant component and all components have size $\Theta(\log n)$. This regime is known as the *sub-critical* regime.

We consider discovery of a random graph in the super-critical regime ($c > 1$). This is the regime of interest, since most real-world networks are well connected rather than having large number of extremely small components. Moreover, the presence of a giant component ensures that the topology can be discovered even with a small fraction of random participants. This is because the participants will most likely belong the giant component, and can thus exchange messages between each other to discover the unknown topology. We limit ourselves to the topology discovery of the giant component in the random graph, and denote the giant component as G_n , unless otherwise mentioned.

2.2 Participation Model

For the given unknown graph topology $G_n = (W_n, E_n)$ over $W_n = \{1, \dots, n\}$ nodes, let $V_n \subset W_n$ be the set of participating nodes which exchange messages amongst each other by routing them along the graph. Let $\rho_n := \frac{|V_n|}{n}$ denote the fraction of participating nodes. It is desirable to have small ρ_n and still reconstruct the unknown topology. We assume that the nodes decide to participate uniformly at random. This ensures that information about all parts of the graph can be obtained, thereby making graph reconstruction feasible. We consider the regime, where $|V_n| = n^{1-\epsilon}$, for some $\epsilon > 0$, meaning that extremely small number of nodes participate in discovering the topology.

Let $H_n := W_n \setminus V_n$ be the set of hidden nodes. The hidden nodes only forward the messages without altering them, and do not provide any additional information for topology discovery. The presence of hidden nodes thus needs to be inferred, as part of our goal of discovering the unknown graph topology.

2.3 Delay Model

The messages exchanged between the participating nodes experience delays along the links in the route. The participating nodes measure the end-to-end delays² between message transmissions and receptions. We consider the challenging scenario that only this end-to-end delay information is available for topology discovery.

Let m be the number of messages exchanged between each pair of participating nodes $i, j \in V_n$. Denote the m samples of end-to-end delays computed from these messages as

$$\mathbf{D}_{i,j}^m := [D_{i,j}(1), D_{i,j}(2), \dots, D_{i,j}(m)]^T.$$

We assume that the routes taken by the m messages are fixed, and we discuss the routing model in the subsequent section. On the other hand, these messages experience different end-to-end delays³ which are drawn identically and independently (i.i.d) from some distribution, described below.

Let D_e denote the random delay along a link $e \in G_n$ (in either direction). We assume that the delays D_{e_1} and D_{e_2} along any two links $e_1, e_2 \in G_n$ are independent. The delays are additive along

²Our algorithms work under any *additive metric* defined on the graph such as link utilization or link loss [42], although the sample complexity, i.e., the number of samples required to accurately estimate the metrics, does indeed depend on the metric under consideration.

³The independence assumption implies that we consider unicast traffic rather than multicast traffic considered in many other works, e.g., in [21].

any route, i.e., the end-to-end delay along a route $\mathcal{P}(i, j)$ between two participants $i, j \in V_n$ is

$$D_{\mathcal{P}(i,j)} := \sum_{e \in \mathcal{P}(i,j)} D_e. \quad (1)$$

We assume that the delay distributions $\{D_e\}_{e \in E_n}$ have bounded variances $\{l(e)\}_{e \in E_n}$ satisfying

$$0 < f \leq l(e) \leq g < \infty, \quad \forall e \in E_n. \quad (2)$$

Further, the family of delay distributions are regular and bounded, as in [12].

The delay distributions $\{D_e\}_{e \in E_n}$ and the graph topology G_n are both unknown, and need to be estimated using messages between participating nodes. We exploit the additivity assumption in (1) to obtain efficient topology discovery algorithms.

2.4 Routing Model

The end-to-end delays between the participating nodes thus depends on the routes taken by the messages. We assume that the messages between any two participants are routed along the shortest path with the lowest number of hops. On the other hand, the nodes cannot select the path with the least delay since the delays along the individual links are unknown and are also different for different messages.

We also consider another scenario, where the participants are able to additionally route messages along the second shortest path. This is a reasonable assumption, since in practice, nodes typically maintain information about the shortest path and an alternative path, should the shortest path fail. The nodes can forward messages along the shortest and the second shortest paths with different headers, so that the destinations can distinguish the two messages and compute the end-to-end delays along the two paths. We will show that this additional information vastly improves the accuracy of topology discovery. These two scenarios are formally defined below.

Scenario 1 (Shortest Path Delays): Each pair of participating nodes $i, j \in V_n$ exchange m messages along the shortest path in G_n , where the shortest path⁴ is with respect to the number of hops. Denote the vector of m end-to-end delays as $\mathbf{D}_{i,j}^m$.

Scenario 2 (Shortest Path and Second Shortest Path Delays): Each pair of participating nodes $i, j \in V_n$ exchange m messages along the shortest path as well as m messages along the second shortest path. The vector of m samples along the second shortest path is denoted by $\tilde{\mathbf{D}}_{i,j}^m$.

3 Reconstruction Guarantees

3.1 Minimal Representation

Our goal is to discover the unknown graph topology using the end-to-end delay information between the participating nodes. However, there can be multiple topologies which explain equally well the end-to-end delays between the participants. This inherent ambiguity in topology discovery with hidden nodes has been previously pointed out in the context of latent tree models [43].

Thus, there is an equivalence class of topologies with different sets of hidden nodes which generate the same end-to-end delay distributions between the participating nodes. We refer to the

⁴If the shortest path between two nodes is not unique, assume that the node pairs randomly pick one of the paths and use it for all the messages.



Figure 1: In minimal representation of a graph, hidden nodes with degree two or lower are merged. See Procedure 1.

Procedure 1 $\tilde{G}_n := \text{Minimal}(G_n; V_n)$ is the minimal representation of G_n given set of participating nodes V_n .

Input: Graph G_n , set of participating nodes V_n , and set of hidden nodes H_n .

Initialize $\tilde{G}_n = G_n$.

while $\exists h \in \tilde{G}_n \cap H_n$ with $\deg(h) \leq 1$ **do**

 Remove h from \tilde{G}_n

end while

Contract all hidden nodes with degree 2 in \tilde{G}_n .

topology with the least number of hidden nodes in this equivalence class as the *minimal representation*. Such a minimal representation does not have redundant hidden nodes. For example, in Fig.1, the graph and its minimal representation are shown. In Procedure 1, we characterize the relationship between a graph and its minimal representation, given a set of participants. The minimal representation is obtained by iteratively removing redundant hidden nodes from the graph.

Thus, any algorithm can only reconstruct the unknown topology up to its minimal representation using only end-to-end delay information between the participating nodes. In sparse random graphs, only a small (but a linear) number of nodes are removed in the minimal representation, and this number decreases with the average degree c . Moreover, given the minimal representation, nodes with degree two in the original graph can be recovered with additional information such as hop counts along the paths between the participants, or can be inferred using the knowledge of the bounds on delay variance in (2). It thus suffices to reconstruct the minimal representation of the original topology, and our goal is to accomplish it using small fraction of participants.

3.2 Performance Measures

We now define performance measures for topology discovery algorithms. It is desirable to have an algorithm which outputs a graph structure which is close to the original graph structure. However, the reconstructed graph cannot be directly compared with the original graph since the hidden nodes introduced in the reconstructed graph are unlabeled and may correspond to different hidden nodes in the original graph. To this end, we require the notion of edit distance defined below.

Definition 1 (Edit Distance) Let F, G be two graphs⁵ with adjacency matrices $\mathbf{A}_F, \mathbf{A}_G$, and let

⁵We consider inexact graph matching where the unlabeled nodes can be unmatched. This is done by adding required number of isolated unlabeled nodes in the other graph, and considering the modified adjacency matrices [15].

V be the set of labeled vertices in both the graphs (with identical labels). Then the edit distance between F, G is defined as

$$\Delta(F, G; V) := \min_{\pi} \|\mathbf{A}_F - \pi(\mathbf{A}_G)\|_1,$$

where π is any permutation on the unlabeled nodes while keeping the labeled nodes fixed.

In other words, the edit distance is the minimum number of entries that are different in \mathbf{A}_F and in any permutation of \mathbf{A}_G over the unlabeled nodes. In our context, the labeled nodes correspond to the participating nodes while the unlabeled nodes correspond to hidden nodes.

Our goal is to output a graph with small edit distance with respect to the minimal representation of the original graph. Ideally, we would like the edit distance to decay as we obtain more delay samples and this is the notion of consistency.

Definition 2 (Consistency) Denote $\widehat{G}_n(\{\mathbf{D}_{i,j}^m\}_{i,j \in V_n})$ as the estimated graph using m delay samples between the participating nodes V_n . A graph estimator $\widehat{G}_n(\{\mathbf{D}_{i,j}^m\}_{i,j \in V_n})$ is structurally consistent if it asymptotically recovers the minimal representation of the unknown topology, i.e.,

$$\lim_{m \rightarrow \infty} \mathbb{P}[\Delta(\widehat{G}_n(\{\mathbf{D}_{i,j}^m\}_{i,j \in V_n}), \widetilde{G}_n; V_n) > 0] = 0. \quad (3)$$

The above definition assumes that the network size n is fixed while the number of samples m goes to infinity. A more challenging setting where both the network size and the number of samples grow is known as the setting of *high-dimensional inference* [52]. In this setting, we are interested in estimating large network structures using a small number of delay samples. We will consider this setting for topology discovery in this paper. Indeed in practice, we have large network structures but can obtain only few end-to-end delay samples with respect to the size of the network. This is formalized using the notion of sample complexity defined below for our setting.

Definition 3 (Sample Complexity) If the number of samples is $m = \Omega(f(n))$, for some function f , such that the estimator $\widehat{G}_n(\{\mathbf{D}_{i,j}^m\}_{i,j \in V_n})$ satisfies

$$\lim_{\substack{m, n \rightarrow \infty \\ m = \Omega(f(n))}} \mathbb{P}[\Delta(\widehat{G}_n(\{\mathbf{D}_{i,j}^m\}_{i,j \in V_n}), \widetilde{G}_n; V_n) = \Omega(g(n))] = 0,$$

for some function $g(n)$, then the estimator \widehat{G}_n is said to have sample complexity of $\Omega(f(n))$ for achieving edit distance of $O(g(n))$.

Thus, our goal is to discover topology in high-dimensional regime, and design a graph estimator that requires a small number of delay samples, and output a graph with a small edit distance.

4 Preliminaries

We now discuss some simple concepts which will be incorporated into our topology discovery algorithms.




Figure 2: Quartet $Q(ab|uv)$. See (6) and (7).

4.1 Delay Variance Estimation

In our setting, topology discovery is based on the end-to-end delays between the participating nodes. Recall that in Section 2.3, we assume general delay distributions on the edges with bounded variances. Our topology discovery algorithms will be based solely on the estimated variances using the end-to-end delay samples.

We use the standard unbiased estimator for variances [36].

$$\hat{l}^m(i, j) := \frac{1}{m-1} \sum_{k=1}^m (D_{i,j}(k) - \bar{D}_{i,j}^m)^2, \quad (4)$$

where $\bar{D}_{i,j}^m$ is the sample mean delay

$$\bar{D}_{i,j}^m := \frac{1}{m} \sum_{k=1}^m D_{i,j}(k). \quad (5)$$

Note that we do not use an estimator specifically tailored for a parametric delay distribution, and hence, the above estimator yields unbiased estimates for any delay distribution.

Our proposed algorithms for topology discovery require only the estimated delay variances $\{\hat{l}(i, j)\}_{i,j \in V}$ as inputs. Indeed, more information is available in the delay samples \mathbf{D}^m . For instance, in [12], the higher-order moments of the delay distribution are estimated using the delay samples and this provides an estimate for the delay distribution. However, we see that for our goal of topology discovery, the estimated end-to-end delay variances suffice and yield good performance.

Recall that $\{l(i, j)\}_{i,j \in V}$ denote the true end-to-end delay variances and that in (1), we assume that the variances are additive along any path in the graph. Along with other assumptions, this implies that the delay variances on the edges form a *metric* on the graph. We will henceforth refer to the variances as “distances” between the nodes and the estimated variances as “estimated distances”. This abstraction also implies that our algorithms will work under input of estimates of any additive metrics.

4.2 Quartet Tests

We first recap the so-called quartet tests, which are building blocks of many algorithms for discovering phylogenetic-tree topologies with hidden nodes [10, 23, 32, 43]. The definition of a quartet is given below. See Fig.2.

Definition 4 (Quartet Test) *Given pairwise distance estimates $\{\hat{l}(i, j)\}_{i,j \in V}$ between the participating nodes $i, j \in V$, A four tuple a, b, u, v is known as a quartet and denoted by $Q = (ab|uv)$ if*

it satisfies the so-called four-point condition

$$\widehat{l}(a, b) + \widehat{l}(u, v) < \min(\widehat{l}(a, u) + \widehat{l}(b, v), \widehat{l}(b, u) + \widehat{l}(a, v)). \quad (6)$$

Thus, the quartet test is local test between tuples of four nodes. For the quartet $Q = (ab|uv)$, let e denote the middle edge of the quartet, i.e., the edge which joins a vertex on the shortest path between a and b to a vertex on the shortest path between u and v (Note that the edge can have zero length if the hidden nodes connecting a, b and u, v are the same.). The estimated length of the middle edge (h_1, h_2) between hidden nodes h_1 and h_2 is given by

$$2\widehat{l}(h_1, h_2) = \widehat{l}(a, u) + \widehat{l}(b, v) - \widehat{l}(a, b) - \widehat{l}(u, v). \quad (7)$$

Similarly, all other edge lengths of the quartet can be calculated through the set of linear equations.

Many phylogenetic-tree reconstruction algorithms proceed by iteratively merging quartets to obtain a tree topology. We employ the quartet test for random graph discovery but the analysis incorporates the presence of cycles. Moreover, we introduce modifications under scenario 2, as outlined in Section 5.2, where second shortest path distances are available in addition to the shortest path distances between the participating nodes.

5 Proposed Algorithms

5.1 Scenario 1

We propose the algorithm RGD1 for discovering random graphs under scenario 1, as outlined in Section 2.4, where only shortest path distance estimates are available between the participating nodes. The idea behind RGD1 is similar to the classical phylogenetic-tree reconstruction algorithms based on quartet tests [23, 43]. However, the effect of cycles on such tests needs to be analyzed, and is carried out in Section 6.1. The algorithm is summarized in Algorithm 2.

The algorithm recursively runs the quartet tests over the set of participating nodes. When quartets are discovered, it merges them with the previously constructed graph using procedure MergeQuartet. This procedure splits the paths between the quartet nodes to introduce new hidden nodes, if such paths exist in the previously constructed graph; otherwise new edges are added. Any edges smaller than a threshold ϵ are contracted, for some chosen constant $\epsilon < f$, where f is the lower bound on the edge lengths of the original graph. It can be shown that the average runtime for RGD1 algorithm is $\tilde{O}(nc^{Rg/2f})$.

The algorithm limits to testing only “short quartets” between nearby participating nodes. Intuitively, this is done to avoid testing quartets on short cycles, since in such scenarios, the quartet tests may fail to reconstruct the graph accurately. Since the random graphs are locally-tree like and contain a small number of short cycles, limiting to short quartets enables us to avoid most of the cycles. The idea of short quartets has been used before (e.g. in [23]) but for a different goal of obtaining low sample complexity algorithm for phylogenetic-tree reconstruction. We carry out a detailed analysis on the effect of cycles on quartet tests in Section 6.1.

In algorithm RGD1, we consider short quartets, where all the estimated distances between the quartet end points are at most $Rg + \tau$, where g is the upper bound on the (exact) edge lengths in the original graph, as assumed in (2). Thus, $R' := Rg/f$ is the maximum number of hops between the end points of a short quartet, where f is the lower bound on the edge lengths. We refer to R' as the *diameter of the quartet*. This needs to be chosen carefully to balance the following two

Algorithm 2 RGD1($\{\widehat{l}(i, j)\}_{i, j \in V_n}; R, g, \tau, \epsilon$) for Topology Discovery Using Shortest-Path Distance Estimates.

Input: Distance estimates between the participating nodes $\{\widehat{l}(i, j)\}_{i, j \in V_n}$, upper bound g on exact edge lengths and parameters $R, \tau, \epsilon > 0$.

Initialize $\widehat{G}_n = (V_n, \emptyset)$.

for $a, b, u, v \in V_n$ such that $\max_{i, j \in \{a, b, u, v\}} \widehat{l}(i, j) < Rg + \tau$ **do**

if $Q(ab|uv)$ (i.e., Quartet) **then**

 Update $\widehat{G}_n \leftarrow \text{MergeQuartet}(\widehat{G}_n, Q(ab|uv); \epsilon)$.

end if

end for

Algorithm 3 RGD2($\{\widehat{l}(i, j), \widehat{l}_2(i, j)\}_{i, j \in V_n}; R, g, \tau, \epsilon$) for Topology Discovery Using Shortest-Path and Second Shortest-Path Distance Estimates.

Input: Shortest-path and second shortest-path distance estimates $\{\widehat{l}(i, j), \widehat{l}_2(i, j)\}_{i, j \in V_n}$, upper bound g on exact edge lengths and parameters $R, \tau, \epsilon > 0$.

Initialize $\widehat{G}_n = (V_n, \emptyset)$.

for $a, b, u, v \in V_n$ **do**

$\mathcal{Q} \leftarrow \{\text{Quartets over } a, b, u, v \text{ using shortest or second shortest distances, which are all less than } Rg + \tau\}$.

 If a quartet occurs twice in \mathcal{Q} , only retain the one with shortest estimated middle edge.

 Update $\widehat{G}_n \leftarrow \text{MergeQuartet}(\widehat{G}_n, \mathcal{Q}; \epsilon)$.

end for

events: encountering short cycles and ensuring that most hidden edges (with at least one hidden end point) are part of short quartets. The parameter τ is chosen to relax the bound, since we have distance estimates, computed using samples, rather than exact distances between the participating nodes.

5.2 Scenario 2

We now consider scenario 2, as outlined in Section 2.4, where second shortest path distance estimates are available in addition to shortest path distance estimates between the participating nodes. We propose RGD2 algorithm for this case, which is summarized in Algorithm 3.

The algorithm RGD2 is an extension of RGD1, where we use the second shortest distances in the quartet tests, in addition to the shortest distances. For each tuple of participating nodes $a, b, u, v \in V_n$, the quartet test in (6) is carried out for all possible combinations of shortest and second shortest distances; only short quartets are retained, where all the distances used for quartet test are less than the specified threshold (which is the same as in RGD1). If the same quartet is formed using different combinations of shortest and second shortest distances, only the quartet with the shorter middle edge, computed using (7), is retained. We clarify the reason behind this rule and give examples on when this can occur in Section 6.1. As before, all these quartets are merged with previously constructed graph using procedure `MergeQuartet`, but with a minor difference that the path lengths need to be checked since there may not be multiple paths between participating nodes with different lengths. The performance analysis for RGD2 is carried out in Section 6.3.

Procedure 4 MergeQuartet($\widehat{G}_n, Q(ab|uv); \epsilon$) for merging a new quartet with current structure \widehat{G}_n .

Input: Current graph \widehat{G} and new quartet $Q(ab|cd)$ with hidden nodes h_1, h_2 (See Fig.2) and threshold ϵ for contracting short edges.

Compute all distances $\{\widehat{l}(e)\}$ inside the quartet.

if \exists path between a, b in \widehat{G} with length $\widehat{l}(a, b; Q)$ **then**

Split path and introduce a new hidden node h'_1 such that $\widehat{l}(a, h'_1; \widehat{G}) = \widehat{l}(a, h_1; Q)$.

else

Introduce node h'_1 and connect it with a and b in \widehat{G} . Assign edge lengths $\widehat{l}(a, h'_1; \widehat{G}) = \widehat{l}(a, h_1; Q)$, $\widehat{l}(b, h'_1; \widehat{G}) = \widehat{l}(b, h_1; Q)$.

end if

Repeat for u, v and introduce hidden node h'_2 .

Join h'_1 and h'_2 and assign it length $\widehat{l}(h'_1, h'_2; \widehat{G}) = \widehat{l}(h_1, h_2; Q)$.

Contract any edges (with at least one hidden end point) if length $< \epsilon$.




Figure 3: A bad quartet: $l(a, b) < l(a, h_1) + l(b, h_1)$.

6 Analysis Under Exact Distances

We now undertake performance analysis for the proposed topology discovery algorithms RGD1 and RGD2. In this section, for simplicity, we first analyze the performance assuming that exact distances between the participating nodes are input to the algorithms. Analysis when distance estimates are input to the algorithms is considered in Section 7.

6.1 Effect of Cycles on Quartet Tests

We now analyze the effect of cycles on quartet tests. Recall that the quartet test is the inequality test in (6), and if this inequality test is satisfied, internal edge lengths of the quartet are computed, and they are added to the output using procedure MergeQuartet. We analyze the effect of cycles on both these aspects of the quartet test.

It is easy to see that the outcome of the quartet test can be incorrect only when the shortest path distances between the end points of the quartet are not the same as the distances along the paths in the quartet. This occurs only when the middle edge of the quartet is part of a (generalized) cycle of length less than $2R'$, where $R' := Rg/f$ is the maximum number of hops between the endpoints of a short quartet. See Fig.3. We refer to such quartets as “bad quartets”.

There are two possible outcomes for bad quartets: (a) the inequality quartet test in (6) fails or (b) the test succeeds. In the former case, the quartet is not reconstructed, while in the latter case the quartet is reconstructed, but with wrong internal edge lengths. See Fig.4 for examples of both the cases. Thus, the performance of RGD1 algorithm depends on the number of bad quartets, and we analyze it in the next section.




Figure 4: Two cases for bad quartets in (a) and (b) with all unit-length edges. In (a), the test in (6) fails, while in (b), the test succeeds but leads to wrong edge estimates, as shown in (c).

For the RGD2 algorithm where second shortest path distances are additionally available, bad quartets do not adversely affect performance. In this scenario, a quartet is correctly recognized when the distances along the quartet correspond to either the shortest or the second shortest path distances between the quartet endpoints. This is because in this case, there is some combination of shortest and second shortest path distances which accurately reconstructs the quartet. Moreover, the availability of the second shortest path distance rules out reconstruction of wrong quartet with incorrect edge lengths since the correct quartet is the one with the shorter middle edge. See Fig.4b for an example of this scenario.

Thus, a quartet is correctly reconstructed if it only consists of shortest or second shortest path distances, and this occurs when the middle edge of the quartet has at most one (generalized) cycle of length less than $2R'$, where R' is the quartet diameter. In Appendix A, we show that there exist no overlapping short cycles, while there a small number of short cycles and hence, we obtain much better accuracy with RGD2 algorithm, compared to RGD1 algorithm.

6.2 Analysis of RGD1

We now provide edit distance guarantees for RGD1 under appropriate choice of maximum quartet diameter $R' := Rg/f$. We analyze the edit distance by counting the number of hidden edges (with at least one hidden end point) which are not recovered correctly under RGD1. A hidden edge is not recovered when one of the following two events occur: (a) it is not part of a short quartet (b) it is part of a bad short quartet. A large value of the quartet diameter R' decreases the likelihood of event (a), while it increases the likelihood of event (b), i.e., we are likely to encounter more cycles as R' is increased. For a fixed value of R' , we analyze the likelihood of these two events and obtain the bound on edit distance stated below.

Assume that the algorithm RGD1 chooses parameter R as $R = \frac{\gamma \log n}{\log c}$, where γ is chosen such that

$$R \leq \frac{f}{g} \left(2 + 2 \frac{\log(n^{0.75}/3)}{\log c} \right) \quad (8)$$

holds, where f and g are lower and upper bounds on exact edge lengths, as assumed in (2). Let the fraction of participating nodes be $\rho_n = n^{-\beta}$, such that

$$\rho c^{\frac{R}{2}} = \omega(1), \quad (9)$$

meaning that $\gamma > 2\beta$. We have the following result.

Theorem 1 (Edit Distance Under RGD1) *The algorithm RGD1 recovers the minimal representation \tilde{G}_n of the giant component of a.e. graph $G_n \sim \mathcal{G}(n, c/n)$ with edit distance*

$$\Delta(\hat{G}_n, \tilde{G}_n; V_n) = \tilde{O}(n^{4\gamma g/f - 4\beta}). \quad (10)$$

Remarks:

(i) The dominant event leading to the edit-distance bound above is the presence of bad quartets due to short cycles in the random graph. In subsequent section, we show that RGD2 algorithm effectively handles this event using the second shortest path distances.

(ii) The parameter R chosen by RGD1 depends on the fraction of participating nodes ρ in (9). Intuitively, R should be large enough, to ensure that most edges lie on short quartets, and can thus be recovered. But larger the R , higher is the resulting edit distance in (10), since we encounter more cycles as we increase R . Thus, the performance improves with the participation fraction, as expected.

(iii) When the ratio of the bounds on the edge lengths g/f is small (i.e., the edge lengths are nearly homogeneous), the edit-distance guarantee in (10) improves, for a fixed R and ρ . This is because we can control the hop lengths of the selected quartets more effectively in this case.

Proof Ideas:

The proof is based on the error events that can cause the quartet tests to fail. The first error event is that an edge which does not occur as a middle edge of a short quartet, meaning that there are not enough participating nodes within distance $R/2$ from it. The second error event is that an edge occurs as a middle edge of a bad quartet, meaning that it is close to a short cycle. We analyze the probability of these events and the resulting edit distance due to these events.

6.3 Analysis of RGD2

We now provide edit distance guarantees for RGD2 algorithm. The analysis is on the lines of the previous section, but we instead analyze the presence of overlapping cycles, as noted in Section 6.1. There are no overlapping short cycles in a random graph, and thus, we can provide a much stronger reconstruction guarantee for the RGD2 algorithm, compared to the RGD1 algorithm.

As before, assume that the parameter R for short quartet as $R = \frac{\gamma \log n}{\log c}$, where γ is chosen such that (8) holds. Let the fraction of participating nodes be $\rho_n = n^{-\beta}$ and (9) holds. We have the following result.

Theorem 2 (Edit Distance Under RGD2) *Under the above conditions, the algorithm RGD2 recovers the minimal representation \tilde{G}_n of the giant component of a.e. graph $G_n \sim \mathcal{G}(n, c/n)$ with edit distance*

$$\Delta(\hat{G}_n, \tilde{G}_n; V_n) = \tilde{O}(n^{6\gamma g/f - 4\beta - 1}). \quad (11)$$

The above result immediately implies that consistent recovery of the minimal representation is possible when there are enough number of participating nodes. We state the result formally below.

Corollary 1 (Consistency Under RGD2) *The algorithm RGD2 consistently recovers the minimal representation \tilde{G}_n of the giant component of a.e. graph $G_n \sim \mathcal{G}(n, c/n)$, when the parameter R and the fraction of participating nodes ρ satisfy*

$$c \frac{6Rg}{f} \rho^4 = o(n), \quad c \frac{R}{2} \rho = \omega(1). \quad (12)$$

Remarks:

(i) From the above constraints, we see that consistent topology recovery is feasible with $\rho = n^{-\beta}$, when $\beta < 12g/f - 4$. Thus, for equal edge lengths ($f = g$), when the number of participants is more than $n^{7/8}$, RGD2 consistently recovers the topology. Intuitively, this is because, with enough participating nodes, we can limit to short distances, and avoid any overlapping cycles, which are the sources of error for RGD2 algorithm.

(ii) Thus, the availability of second shortest distances makes consistent topology discovery possible with a sub-linear number of participating nodes, while consistent recovery is not tractable under RGD1 using only shortest-path distances between a sub-linear number of participants.

(iii) Comparing the edit-distance guarantees for RGD1 and RGD2 algorithms in Theorem 1 and Theorem 2, we note that RGD2 reconstructs the topology with better accuracy than RGD1, when $\beta > \frac{f}{4g}$, for the same fixed parameter R . Thus, RGD2 has better performance than RGD1, when there are at least $n^{0.75}$ number of participants for the homogeneous case ($f = g$).

Proof Ideas:

The proof is on similar lines as in Theorem 1, but with modified error events that cause the quartet tests to fail. As before, the first error event is that an edge which does not occur as a middle edge of a short quartet. The second error event is now that an edge is close to two overlapping short cycles instead of being close to a single short cycle. This event does not occur in random graphs for sufficiently short lengths, and thus, we see a drastic improvement in edit distance.

7 Analysis Under Samples

We have so far analyzed the performance of RGD1 and RGD2 algorithms when exact distances (i.e., delay variances) are input to the algorithm. We now analyze the scenario when instead only delay samples are available and estimated variances are input to the algorithm.

We show that the proposed algorithms have low sample complexity, meaning they require slow scaling of number of samples compared to the network size to achieved guaranteed performance. The result is given below.

Theorem 3 (Sample Complexity) *The edit distance guarantees under RGD1 and RGD2 algorithms, as stated in Theorem 1 and Theorem 2, are achieved under input of estimated delay variances, if the number of delay samples satisfies*

$$m = \Omega(\text{poly}(\log n)). \tag{13}$$

Thus, the sample complexity of RGD1 and RGD2 algorithms is $\text{poly}(\log n)$. In other words, the size of the network n can grow much faster than the number of delay samples m , and we can still obtain good estimates of the network. This implies with $m = \Omega(\text{poly}(\log n))$ samples, we can consistently discover the topology under RGD2 algorithm, given sufficient fraction of participating nodes.

Proof Ideas:

The proof follows from Azuma-Hoeffding inequality for concentration of individual variance estimates, as in [12, Proposition 1], and then consider the union bound over various events.

8 Converse Results & Discussion

8.1 Fraction of Participating Nodes

We have so far provided edit distance guarantees for the proposed topology discovery algorithms. In this section, we provide a lower bound on the fraction of participating nodes required for any algorithm to recover the original graph up to a certain edit distance guarantee.

We can obtain a meaningful lower bound only when the specified edit distance is lower than the edit distance between a given graph and an independent realizations of the random graph. Otherwise, the edit distance guarantee can always be realized by a random construction of the output graph. To this end, we first prove a lower bound on the edit distance between two independent realizations of the random graph.

Let $\mathcal{D}(G; \delta)$ denote the set of all graphs which have edit distance of at most δ from G

$$\mathcal{D}(G; \delta) := \{F : \Delta(F, G; \emptyset) > \delta\}. \quad (14)$$

Lemma 1 (Lower Bound on Edit Distance) *Almost every random graph $G_n \sim \mathcal{G}(n, c/n)$ has an edit distance at least $(0.5c - 1)n$ from any given graph F_n .*

Proof: First, we have for any graph F_n

$$|\mathcal{D}(F_n; \delta n)| \leq n! \cdot \binom{\frac{n^2}{2}}{\delta n} < n^{(\delta+1)n} 3^{\delta n}, \quad (15)$$

since we can permute the n vertices and change at most δn entries in the adjacency matrix \mathbf{A}_F and we use the bound that $\binom{N}{k} \leq \frac{N^k}{k!} \leq \left(\frac{N}{k}\right)^k 3^k$. Let \mathcal{B} denote the set of graphs having exactly $\frac{cn}{2}$ edges and the size of \mathcal{B} is

$$|\mathcal{B}| = \binom{\frac{n^2}{2}}{\frac{cn}{2}} \geq \left(\frac{n^2}{cn}\right)^{cn/2} = \left(\frac{n}{c}\right)^{cn/2}.$$

We can now bound the probability that a random graph $G_n \sim \mathcal{G}(n, c/n)$ belongs to set $\mathcal{D}(F_n; \delta n)$ for any given graph F_n is

$$\begin{aligned} \mathbb{P}[G_n \in \mathcal{D}(F_n; \delta n)] &\leq \frac{\mathbb{P}[G_n \in \mathcal{D}(F_n; \delta n)]}{\mathbb{P}[G_n \in \mathcal{B}]} \\ &\leq \frac{|\mathcal{D}(F_n; \delta n)| \max_{g \in \mathcal{D}(F_n; \delta n)} \mathbb{P}[G_n = g]}{|\mathcal{B}| \min_{g \in \mathcal{B}} \mathbb{P}[G_n = g]} \\ &\stackrel{(a)}{\leq} \frac{|\mathcal{D}(F_n; \delta n)|}{|\mathcal{B}|} = n^{(\delta+1-c/2)n} 3^{\delta n}, \end{aligned}$$

where inequality (a) is due to the fact that $\min_{g \in \mathcal{B}} \mathbb{P}[G_n = g] \geq \max_{g \in \mathcal{S}(F_n)} \mathbb{P}[G_n = g]$ (i.e., the mode of the binomial distribution). Hence, $\mathbb{P}[G_n \in \mathcal{D}(F_n; \delta n)]$ decays to zero as $n \rightarrow \infty$, when $\delta < 0.5c - 1$. \square

Thus, for any given graph, a random graph does not have edit distance $< (0.5c - 1)n$ from it. It is thus reasonable to expect for any graph reconstruction algorithm to achieve an edit distance less than $(0.5c - 1)n$, since otherwise, a random choice of the output graph could achieve the same edit distance. We now provide a lower bound on the fraction of the participating nodes such that no algorithm can reconstruct the original graph up to an edit distance less than $(0.5c - 1)n$.



Figure 5: Example of two graphs with unit lengths where nodes a, b, c, d are participating. Even under all path length information between the participating nodes, the two graphs cannot be distinguished.

Theorem 4 (Lower Bound) For $G_n \sim \mathcal{G}(n, c/n)$ and any set of participants V_n , for any graph estimator \hat{G}_n using (exact) shortest path distances between the participating node pairs, we have

$$\mathbb{P}[\Delta(\hat{G}_n, G_n; V) > \delta n] \rightarrow 1, \text{ when}$$

$$|V|^2 < Mn(0.5c - \delta - 1) \frac{\log n}{\log \log n}, \quad (16)$$

for a small enough constant $M > 0$ and any $\delta < (0.5c - 1)$.

Thus, no algorithm can reconstruct G_n up to edit distance δn , for $\delta < 0.5c - 1$, if the number of participating nodes is below a certain threshold. From Lemma 1, almost every random graph has an edit distance greater than $(0.5c - 1)n$ from a given graph. Thus, when the number of participating nodes is below a certain threshold, accurate reconstruction by any algorithm is impossible.

Remarks:

- (i) The lower bound does not require that the participating nodes are chosen uniformly and holds for any set of participating nodes of given cardinality.
- (ii) The lower bound is analogous to a strong converse in information theory [19] since it says that the probability of edit distance being more a certain quantity goes to one (not just bounded away from zero).
- (iii) The result is valid even for the scenario where second shortest path distances are used since the maximum second shortest path distance is also $O(\log n)$.
- (iv) We have earlier shown that our algorithms RGD1 and RGD2 have good performance under a sub-linear number of participants. Closing the gaps in the exponents between lower bound and achievability is of interest.

Proof Ideas:

The proof is based on information-theoretic covering type argument, where cover the range of the estimator with random graphs of high likelihood. Using bounds on binomial distribution, we obtain the desired lower bound.

8.2 Non-Identifiability of General Topologies

Our proposed algorithms require the knowledge of shortest and second shortest path distances. Performance analysis reveals that the knowledge of second shortest path can greatly improve the

accuracy of topology discovery for random graphs. We now address the question if this can be accomplished in general.

To this end, we provide a counter-example in Fig.5, where a significant fraction of nodes are participating, and we are given distances along all the paths between the participants; yet, the topology cannot be correctly identified by any algorithm. This reveals a fundamental non-identifiability of general topologies using only a subset of participating nodes.

8.3 Relationship to Phylogenetic Trees

We note some key differences between the phylogenetic-tree model [22] and the additive delay model employed in this paper. In phylogenetic trees, sequences of extant species are available, and the unknown phylogenetic tree is to be inferred from these sequences. The phylogenetic-tree models the series of mutations occurring as the tree progresses and new species are formed. Efficient algorithms with low sample complexity have been proposed for phylogenetic-tree reconstruction, e.g., in [20,23].

In the phylogenetic-tree model, the correlations along the phylogenetic tree decay exponentially with the number of hops. This implies that long-range correlations (between nodes which are far away) are “hard” to estimate, and require large number of samples (compared to the size of the tree) to find an accurate estimate. However, under the delay model, the delays are additive along the edges, and even long-range delays can be shown to be “easy” to estimate. Hence, the delay model does not require the more sophisticated techniques developed for phylogenetic-tree reconstruction (e.g., [20]), in order to achieve low sample complexity. However, the presence of cycles complicates the analysis for delay-based reconstruction of random graphs. Moreover, we developed algorithms when additional information is available in the form of second shortest-path distances. Such information cannot be obtained from phylogenetic data. We demonstrated that this additional information leads to drastic improvement in the accuracy of random-graph discovery.

9 Conclusion

In this paper, we considered discovery of sparse random graph topologies using a sub-linear number of uniformly selected participants. We proposed local quartet-based algorithms which exploit the locally tree-like property of sparse random graphs. We first showed that a sub-linear edit-distance guarantee can be obtained using end-to-end measurements along the shortest paths between a sub-linear number of participants. We then considered the scenario where additionally, second shortest-path measurements are available, and showed that consistent topology recovery is feasible using only a sub-linear number of participants. Finally, we establish a lower bound on the edit distance achieved by any algorithm for a given number of participants. Our algorithms are simple to implement, computationally efficient and have low sample complexity.

There are many interesting directions to explore. Our algorithms require the knowledge of the bounds on the delay variances (i.e., edge lengths), and algorithms which remove these requirements can be explored. Our algorithms are applicable for other locally tree-like graphs as well, while the actual performance indeed depends on the model employed. Exploring how the reconstruction performance changes with the random graph model considered is of interest. Another scenario is when instead of an additive-delay model considered here, the links could represent correlations or similarities between the nodes, as in the phylogenetic-tree model [22]. Topology discovery of such

models under a random graph structure is a challenging problem.

Acknowledgements

The authors thank Venkat Chandrasekharan at MIT for extensive discussions and helpful comments.

A Properties of Random Graphs

We now recap the bounds on neighborhood size of random graph from [46, Lemma 2.4]. Let $0 < \delta < 1/1000$ be an arbitrarily small constant. A vertex v is said to have a *large regular neighborhood* if either: $|\Gamma_t(v)| < M$ for all t , or setting $t^- := \min\{t : |\Gamma_t(v)| \geq M\}$ and $t^+ := t^- + \log(n^{0.75}/M)/\log c$, we have

$$(1 - \delta)c^{t-t^-+1}|\Gamma_{t-t^-}(v)| \leq |\Gamma_t(v)| \leq (1 + \delta)c^{t-t^-+1}|\Gamma_{t-t^-}(v)|,$$

for all $t^- \leq t \leq t^+$.

Lemma 2 (Bounds on Neighborhood) *The probability that a vertex does not have a large regular neighborhood is $o(n^{-1})$.*

Lemma 3 (Cycles in Random Graphs) *In $G_n \sim \mathcal{G}(n, \frac{c}{n})$, the expected number of cycles of lengths l is $O(c^l)$. Moreover, the number of two overlapping cycles of length l , denoted by H_l , satisfies*

$$\mathbb{E}[N_{H_l}] = O(n^{-1}c^{2l+1}). \tag{17}$$

Thus, there are a.a.s. no overlapping cycles of length less than $\frac{(1-\delta)\log n}{2\log c}$ for $\delta > 0$.

Proof: The proof is along the lines of [13, Cor. 4.9], but we specialize it for cycles. By counting argument, the expected number of cycles is given by

$$\mathbb{E}[N_{C(l)}] = \binom{n}{l} \frac{l!}{2l} \left(\frac{c}{n}\right)^l = O(c^l).$$

Let number of vertices in H , be $|v(H_l)| = s$ with $l \leq s \leq 2l$. Note that the number of edges $|H_l| \geq s + 1$ to be overlapping cycles. Hence,

$$\begin{aligned} \mathbb{E}[N_{H_l}] &\leq \binom{n}{s} (s!) \left(\frac{c}{n}\right)^{s+1} \\ &= O(n^{-1}c^{s+1}), \end{aligned}$$

and we obtain the desired result. □

B Proof of Theorem 1

Let \tilde{G}_n be the minimal representation of the original graph G_n , let \hat{G}_n be the graph output under RGD1 (with exact distances). We can express the edit distance in terms of the following events: let $\mathcal{E}_1(e; \tilde{G}_n, V_n)$ denote the event that the edge e is not a middle edge in any short quartet. Let $\mathcal{E}_2(v; \tilde{G}_n, V_n)$ denote the event that the node v is the middle node of a bad short quartet, and let

K_v denote the number of such bad short quartets (with participating nodes as end points). The edit distance satisfies

$$\Delta(\widehat{G}_n, \widetilde{G}_n; V_n) \leq \sum_{v \in \widetilde{G}_n} \left\{ (\deg(v; \widetilde{G}_n) + 3K_v) \mathbb{I}[\mathcal{E}_2(v)] + \sum_{e: i \subset e} \mathbb{I}[\mathcal{E}_1(e)] \right\}. \quad (18)$$

This is because under event $\mathcal{E}_1(e)$, no short quartet contains e as a middle edge and hence, it is not reconstructed. When v is the middle node of a bad quarter, either is not reconstructed, in which case, this contributes an edit distance of at most $\deg(v)$ or the quartet is reconstructed with wrong edge lengths. In the latter case, this amounts to adding a spurious hidden node with degree 3, and we can add a maximum of $3K_v$ such edges, where K_v is the number of bad short quartets. Moreover, from the definition of minimal representation \widetilde{G}_n in Procedure 1, we have for $v \in \widetilde{G}_n$,

$$\mathbb{P}[\mathcal{E}_2(v; \widetilde{G}_n, V_n)] \leq \mathbb{P}[\mathcal{E}_2(v; G_n, V_n) | \deg(v) \geq 3],$$

and similarly for $\mathcal{E}_1(e)$, we can relate the event in minimal representation and the original graph. We now analyze events \mathcal{E}_1 and \mathcal{E}_2 separately.

For the event $\mathcal{E}_1(e)$, let v be one of the end points of e . We have

$$\mathbb{P}[\mathcal{E}_1(e; G_n, V_n)] \leq 4\mathbb{P}[V_n \cap B_{R/2}(v; G_n) = \emptyset],$$

and

$$\mathbb{P}[V_n \cap B_{R/2}(v; G_n) = \emptyset | |B_{R/2}(v; G_n)| \geq k] \leq (1 - \rho)^k.$$

We have a lower bound on $|B_{R/2}(v)|$ from Lemma 2, by noting that $M = 3$ and $t^- = 1$, due to conditioning on the node degree. Hence, for

$$R \leq 2 + 2 \log(n^{0.75}/3) / \log c, \quad (19)$$

with probability $1 - o(n^{-1})$, given that $\deg(e) \geq 4$, we have, with probability $1 - o(n^{-1})$,

$$|B_{R/2}(v)| \geq (1 - \delta)c^{R/2},$$

for some small $\delta > 0$. Thus,

$$\mathbb{P}[V \cap B_{R/2}(v) = \emptyset | \deg(v) \geq 3] \leq (1 - \rho)^{(1-\delta)c^{R/2}}. \quad (20)$$

For the second event \mathcal{E}_2 , that the edge v is a part of a bad quartet, this occurs when it is part of a (generalized) cycle of length less than $2R'$,

$$\mathbb{P}[\mathcal{E}_2(v; G_n, V_n)] = \mathbb{P}[v \in \mathcal{C}(2R')]$$

We have from Lemma 3,

$$\mathbb{P}[e \in \mathcal{C}(2R')] = O(n^{-1}c^{2R'}).$$

The number of bad short quartets K_v satisfies

$$K_v = \widetilde{O}(\rho^4 c^{2R'}),$$

since $K_v \leq |V_n \cap B_{R'}(v; G_n)|^4$ and

$$\mathbb{E}[|V_n \cap B_{R'}(v; G_n)|] = \rho c^{R'},$$

and using Chernoff bounds, we have $|V_n \cap B_{R'}(v; G_n)| = \tilde{O}(\rho c^{R'/2})$ with probability $1 - o(n^{-1})$.

Thus, the expected edit distance is

$$\begin{aligned} \mathbb{E}[\Delta(\widehat{G}_n, \widetilde{G}_n; V_n)] &= O(n(1 - \rho)^{(1-\delta)c^{R/2}}) \\ &\quad + \tilde{O}(c^{4R'} \rho^4). \end{aligned} \tag{21}$$

Let $R' = Rg/f = \frac{\gamma g \log n}{f \log c}$, where γ is chosen to satisfy (19). Thus $c^{R'} = n^{\gamma g/f}$. Let $\rho_n = n^{-\beta}$. We have

$$(1 - \rho)^{\Theta(c^{R/2})} = O(\exp[-n^{\gamma/2-\beta}]),$$

when $\gamma > 2\beta$. The second term in (21) is $\tilde{O}(n^{4\gamma g/f-4\beta})$, and is the dominant error event. Thus, the expected edit distance is

$$\mathbb{E}[\Delta(\widehat{G}_n, \widetilde{G}_n; V_n)] = \tilde{O}(n^{4\gamma g/f-4\beta}).$$

By Markov inequality, we have the result. \square

C Proof of Theorem 2

The proof follows the lines of proof of Theorem 1. We again have error events \mathcal{E}_1 and \mathcal{E}_2 which lead to a bound on edit distance (18). As before, let $\mathcal{E}_1(v; \widehat{G}_n, V_n)$ denote the event that the node v is not a middle node in any short quartet and $\mathcal{E}_2(v)$ is the event that the node v is the middle node of a bad short quartet. However, now, the definition of a bad quartet is different: it occurs only when node v part of at least two overlapping generalized cycles of length less than $2R'$, and denote such structures as $H_{2R'}$.

The analysis for \mathcal{E}_1 is same as in proof of Theorem 1. For \mathcal{E}_2 , from Lemma 3, we have,

$$\mathbb{P}[v \in H_{2R'}] = O(n^{-2} c^{4R'}).$$

Thus, the expected edit distance is

$$\begin{aligned} \mathbb{E}[\Delta(\widehat{G}_n, \widetilde{G}_n; V_n)] &= O(n(1 - \rho)^{(1-\delta)c^{R/2}}) \\ &\quad + \tilde{O}(n^{-1} c^{6R'} \rho^4). \end{aligned}$$

For $R = \frac{\gamma \log n}{\log c}$ and $\rho_n = n^{-\beta}$, we have the desired result. \square

D Proof of Theorem 3

The proof follows the lines of sample complexity results in [12]. From [12, Proposition 1], we have concentration bounds for delays (distances) under m samples as

$$\mathbb{P}[|\widehat{l}^m(i, j) - l(i, j)| > \epsilon] \leq 2 \exp\left[-\frac{m\epsilon^2}{MR^3}\right],$$

for any $\epsilon > 0$, some constant $M > 0$, and for all $i, j \in V_n$. Taking union bound over all node pairs, we see that when $m = \Omega(\text{poly}(\log n))$, we have concentration of all the distances and we have the desired result. \square

E Proof of Theorem 4

We use a covering argument for obtaining the lower bound, inspired by [14, Thm. 1]. For reconstructed graph \widehat{G} using shortest path distances between $O(|V(G)|^2/2)$ node pairs, the range $\mathcal{R}(\widehat{G})$ of the estimator is bounded by

$$|\mathcal{R}(\widehat{G})| \leq (\text{Diam}(G))^{|V(G)|^2/2},$$

since the delay variances on the edges are assumed to be known exactly, and the shortest path can range from 1 to $\text{Diam}(G)$. For $G \sim \mathcal{G}(n, c/n)$, the diameter⁶ is $O(\log n)$ w.h.p. Let $\mathcal{S}(\widehat{G}; \delta n)$ denote all the graphs which are within edit distance of δn of the graphs in range $\mathcal{R}(\widehat{G})$

$$\mathcal{S}(\widehat{G}; \delta n) := \{F : \Delta(F, G') \leq \delta n, \text{ for some } G' \in \mathcal{R}(\widehat{G})\}.$$

Thus, using (14) and (15),

$$|\mathcal{S}(\widehat{G}; \delta n)| \leq \bigcup_{G' \in \mathcal{R}(\widehat{G})} |\mathcal{D}(G'; \delta n)| \leq |\mathcal{R}(\widehat{G})| n^{(\delta+1)n} \mathfrak{z}^{\delta n}.$$

For the original graph $G \sim \mathcal{G}(n, c/n)$, we have the required probability

$$\begin{aligned} \mathbb{P}[\Delta(\widehat{G}, G; V) > \delta n] &= \sum_{g \in \mathcal{S}^c} \mathbb{P}[\Delta(\widehat{G}, g; V)] > \delta n] \mathbb{P}(G = g) \\ &\quad + \sum_{g \in \mathcal{S}} \mathbb{P}[\Delta(\widehat{G}, g; V)] > \delta n] \mathbb{P}(G_n = g) \\ &\geq \sum_{g \in \mathcal{S}^c} \mathbb{P}[\Delta(\widehat{G}, g; V) > \delta n] \mathbb{P}(G_n = g) \\ &\stackrel{(a)}{=} \sum_{g \in \mathcal{S}^c} \mathbb{P}(G_n = g) \\ &\stackrel{(b)}{=} 1 - \sum_{g \in \mathcal{S}} \mathbb{P}(G_n = g), \end{aligned} \tag{22}$$

where equality (a) is due to the fact that $\mathbb{P}[\Delta(\widehat{G}, g; V) > \delta n] = 1$ for all $g \in \mathcal{S}^c$ and (b) is due to $\sum_{g \in \mathcal{S}} \mathbb{P}(G = g) + \sum_{g \in \mathcal{S}^c} \mathbb{P}(G = g) = 1$.

Now we provide an asymptotic upper bound for the term $\Upsilon := \sum_{g \in \mathcal{S}} \mathbb{P}(G = g)$. Furthermore, let $e_g \in \{1, \dots, \binom{n}{2}\}$ denote the number of edges in the graph $g \in \mathcal{G}_n$. Then,

$$\mathbb{P}(G = g) = \left(\frac{c}{n}\right)^{e_g} \left(1 - \frac{c}{n}\right)^{\binom{n}{2} - e_g}. \tag{23}$$

We have the general result that for graphs $g_1, g_2 \in \mathcal{G}_n$

$$e_{g_1} \leq e_{g_2} \quad \Rightarrow \quad \mathbb{P}(G = g_1) \geq \mathbb{P}(G = g_2). \tag{24}$$

Define

$$z := \min \left\{ l \in \mathbb{N} : \sum_{k=1}^l \binom{\binom{n}{2}}{k} \geq |\mathcal{S}| \right\} \tag{25}$$

⁶The diameter of $G(n, \frac{c}{n})$ is $C(c) \log n$ [26], where $C(c) = \frac{1}{\log c} + \frac{2}{c} + O(\frac{\log c}{c^2})$ as $c \rightarrow \infty$.

We obtain

$$z \leq \frac{|V|^2 \log \log n}{4 \log n} + n(\delta + 1) + o(1).$$

Thus,

$$\begin{aligned} \Upsilon &:= \sum_{g \in \mathcal{S}} \mathbb{P}(G = g) \\ &\leq \sum_{k=0}^z \binom{\binom{n}{2}}{k} \left(\frac{c}{n}\right)^k \left(1 - \frac{c}{n}\right)^{\binom{n}{2} - k} \\ &\stackrel{(a)}{\leq} \exp \left[-\frac{4}{nc} \left(n(0.5c - \delta - 1) - O\left(\frac{|V|^2 \log \log n}{\log n}\right) - o(1) \right)^2 \right] \end{aligned}$$

where inequality (a) follows from the fact that $\Pr(\text{Bin}(N, q) \leq k) \leq \exp(-\frac{2}{Nq}(Nq - k)^2)$ for $k \leq Nq$ with the identifications $N = \binom{n}{2}$, $q = c/n$ and $k = z$. Finally, we observe from (a) that if $|V|^2 < Mn(0.5c - \delta - 1)\frac{\log n}{\log \log n}$ for small enough $M > 0$, then $\Upsilon \rightarrow 0$ as $n \rightarrow \infty$ and we obtain the required result. \square

References

- [1] Cooperative Analysis for Internet Data Analysis, (CAIDA). <http://www.caida.org/tools/>.
- [2] Internet Mapping Project. <http://www.cheswick.com/ches/map/>.
- [3] mtrace—Print multicast path. <ftp://ftp.parc.xerox.com/pub/net-research/ipmulti>.
- [4] The Skitter Project. <http://www.caida.org/tools/measurement/skitter/>.
- [5] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi. Spread of (mis) information in social networks. *Games and Economic Behavior*, 2010.
- [6] D. Alderson, H. Chang, M. Roughan, S. Uhlig, and W. Willinger. The many facets of internet topology and traffic. *AIMS J. on Networks and Heterogeneous Media*, 1(4):569, 2006.
- [7] N. Alon and J. Spencer. *The probabilistic method*. Wiley-Interscience, 2000.
- [8] A. Anandkumar, C. Bisdikian, and D. Agrawal. Tracking in a Spaghetti Bowl: Monitoring Transactions Using Footprints. In *Proc. of ACM SIGMETRICS*, Annapolis, Maryland, USA, June 2008.
- [9] L. Backstrom, C. Dwork, and J. Kleinberg. Wherefore Art Thou r3579x?: Anonymized Social Networks, Hidden Patterns, and Structural Steganography. In *Proc. of ACM Intl. Conf. on World Wide Web*, Banff, Canada, May 2007.
- [10] H.-J. Bandelth and A. Dress. Reconstructing the shape of a tree from observed dissimilarity data. *Adv. Appl. Math*, 7:309–43, 1986.
- [11] Z. Beerliova, F. Eberhard, T. Erlebach, A. Hall, M. Hoffmann, M. Mihal ak, and L. Ram. Network Discovery and Verification. *IEEE Journal on Selected Areas in Communications*, 24(12):2168, 2006.

- [12] S. Bhamidi, R. Rajagopal, and S. Roch. Network Delay Inference from Additive Metrics. *To appear in Random Structures and Algorithms, Arxiv preprint math/0604367*, 2010.
- [13] B. Bollobás. *Random Graphs*. Academic Press, 1985.
- [14] G. Bresler, E. Mossel, and A. Sly. Reconstruction of Markov Random Fields from Samples: Some Observations and Algorithms. In *Intl. workshop APPROX Approximation, Randomization and Combinatorial Optimization*, pages 343–356. Springer, 2008.
- [15] G. Bunke et al. Inexact graph matching for structural pattern recognition. *Pattern Recognition Letters*, 1(4):245–253, 1983.
- [16] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu. Network Tomography: Recent Developments. *Stat. Science*, 19:499–517, 2004.
- [17] S. Choi and J. Kim. Optimal query complexity bounds for finding graphs. In *Proc. of annual ACM symposium on Theory of computing*, pages 749–758, 2008.
- [18] F. Chung, M. Garrett, R. Graham, and D. Shallcross. Distance realization problems with applications to Internet tomography. *Journal of Computer and System Sciences*, 63(3):432–448, 2001.
- [19] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.
- [20] C. Daskalakis, E. Mossel, and S. Roch. Optimal phylogenetic reconstruction. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 159–168, 2006.
- [21] N. Duffield, J. Horowitz, F. Presti, and D. Towsley. Multicast topology inference from end-to-end measurements. *Advances in Performance Analysis*, 3:207–226, 2000.
- [22] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison. *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge Univ. Press, 1999.
- [23] P. L. Erdős, L. A. Székely, M. A. Steel, and T. J. Warnow. A few logs suffice to build (almost) all trees: Part ii. *Theoretical Computer Science*, 221:153–184, 1999.
- [24] B. Eriksson, P. Barford, R. Nowak, and M. Crovella. Learning Network Structure from Passive Measurements. In *Proc. of the ACM SIGCOMM conference on Internet measurement*, Kyoto, Japan, Aug. 2007.
- [25] T. Erlebach, A. Hall, and M. Mihal'ak. Approximate Discovery of Random Graphs. *Lecture Notes in Computer Science*, 4665:82, 2007.
- [26] D. Fernholz and V. Ramachandran. The diameter of sparse random graphs. *Random Structures and Algorithms*, 31(4):482–516, 2007.
- [27] S. Fortunato. Community Detection in Graphs. *Physics Reports*, 2009.
- [28] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring Networks of Diffusion and Influence. In *Proc. of the ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining*, 2010.

- [29] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery. In *IEEE INFOCOM*, Tel-Aviv, Israel, June 2000.
- [30] M. Gunes and K. Sarac. Resolving anonymous routers in Internet topology measurement studies. In *Proc. of IEEE INFOCOM*, pages 1076–1084, 2008.
- [31] Y. He, G. Siganos, and M. Faloutsos. Internet Topology. In R. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 4930–4947. Springer, 2009.
- [32] T. Jiang, P. E. Kearney, and M. Li. A polynomial-time approximation scheme for inferring evolutionary trees from quartet topologies and its application. *SIAM J. Comput.*, 30(6):1942–1961, 2001.
- [33] M. Jovanović, F. Annexstein, and K. Berman. Modeling peer-to-peer network topologies through small-world models and power laws. In *IX Telecommunications Forum, TELFOR*, 2001.
- [34] S. Kandula, D. Katabi, and J. Vasseur. Shrink: A Tool for Failure Diagnosis in IP Networks. In *Proc. of ACM SIGCOMM Workshop on Mining network data*, Philadelphia, PA, Aug. 2005.
- [35] S. Khuller, B. Raghavachari, and A. Rosenfeld. Landmarks in graphs. *Discrete Appl. Math.*, 70(3):217–229, 1996.
- [36] E. Lehmann. *Theory of Point Estimation*. Chapman & Hall, New York, NY, 1991.
- [37] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting Positive and Negative Links in Online Social Networks. In *ACM WWW Intl. Conf. on World Wide Web*, 2010.
- [38] H. Mazzawi. Optimally Reconstructing Weighted Graphs Using Queries. In *Symposium on Discrete Algorithms*, pages 608–615, 2010.
- [39] A. Motter and Y. Lai. Cascade-based Attacks on Complex Networks. *APS Physical Review E*, 66(6):65–102, 2002.
- [40] S. Myers and J. Leskovec. On the Convexity of Latent Social Network Inference. In *Proc. of NIPS*, 2010.
- [41] M. Newman, D. Watts, and S. Strogatz. Random graph models of social networks. *Proc. of the National Academy of Sciences of the United States of America*, 99(Suppl 1), 2002.
- [42] J. Ni, H. Xie, S. Tatikonda, and Y. Yang. Efficient and dynamic routing topology inference from end-to-end measurements. *Networking, IEEE/ACM Transactions on*, 18(1):123–135, 2010.
- [43] J. Pearl. *Probabilistic Reasoning in Intelligent Systems—Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [44] L. Reyzin and N. Srivastava. Learning and verifying graphs using queries with a focus on edge counting. *Lecture Notes in Computer Science*, 4754:285, 2007.
- [45] L. Reyzin and N. Srivastava. On the longest path algorithm for reconstructing trees from distance matrices. *Information Processing Letters*, 101(3):98–100, 2007.

- [46] O. Riordan and N. Wormald. The diameter of sparse random graphs. *Combinatorics, Probability and Computing*, 19(5-6):835–926, 2010.
- [47] D. Shah and T. Zaman. Detecting Sources of Computer Viruses in Networks: Theory and Experiment. In *Proc. of ACM Sigmetrics*, New York, NY, June 2010.
- [48] Y. Shavitt and E. Shir. DIMES: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5):74, 2005.
- [49] M. Shih and A. Hero. Unicast inference of network link delay distributions from edge measurements. In *Proc. of IEEE ICASSP*, volume 6, pages 3421–3424, 2002.
- [50] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP Topologies with Rocketfuel. *IEEE/ACM Tran. on networking*, 12(1):2–16, 2004.
- [51] Y. Vardi. Network Tomography: Estimating Source-Destination Traffic Intensities from Link Data. *J. of the American Statistical Association*, 91(433), 1996.
- [52] M. Wainwright and M. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- [53] F. Wu, B. Huberman, L. Adamic, and J. Tyler. Information Flow in Social Groups. *Physica A: Statistical and Theoretical Physics*, 337(1-2):327–335, 2004.
- [54] B. Yao, R. Viswanathan, F. Chang, and D. Waddington. Topology inference in the presence of anonymous routers. In *Proc. of IEEE INFOCOM*, 2003.