Metrics Useful in Network Tomography Studies.

Yehuda Vardi

Abstract— To facilitate the development of statistical methods geared at analyzing data from networks, it is important to have metrics which define and measure distances between a network's links, its paths, and also between different networks. This is particularly important in the rapidly growing area of network tomography which plays a central role in studies of data, communication, and internet traffic. We propose such metrics, outline some of their properties, and motivate them with two very recent applications. The proposed metrics are simple yet have appealing properties.

***Index Terms*—distance functions on graphs, network routing, data and computer networks, graph topology, internet exploration.**

1. INTRODUCTION

Networks are complicated mathematical structures. They are abundant around us: social, behavioral, traffic, communication, information, and data networks are some of the examples where network structures and data are common. Thus it is important to develop tools for the statistical analysis of network data. Such tools typically depend on having a metric which quantifies notions of "similarity", "closeness", and facilitates ordered-comparisons. The metric is a ground-level component of any statistical study, as it determines the statistical properties (robustness, efficiency, etc.) of the analysis. With this in mind, we propose metrics for measuring distances between links, between paths, and between networks. We focus on data-networks with fixed routing schemes, as defined and discussed in Vardi (1996). Coates et al (2002) noted that this model is useful for a broad range of problems in this rapidly growing area. We take this as a justification to focus first on such networks. Extensions, challenges, motivating examples, and potential applications in network tomography are noted.

2. METHODOLOGY

Consider a strongly connected directed network, where there exists a directed path between any pair of nodes in the network. An edge between two immediately connected nodes is referred to as a "link", L(i), and is indexed by i= $(i_1 \rightarrow i_2)$, i=1,…,I. A set of connected links leading from a source-node j_1 to a destination-node j_2 is referred to as a "path" or "trail", T(j), with j = (j_1,j_2), j= 1,…,J, indexing the source-destination (SD) pairs.

Consider now a fixed-routing protocol: for each SD pair, j, the same unique path, T(j), always carries the traffic from j_1 to j_2. The routing protocol is described by a "routing matrix", A, with entries $a_{i,j}$=1 if L(i) is in T(j) and $a_{i,j}$=0 if L(i) is not in T(j), i = 1,…,I, j= 1,…,J. We now define distance functions for links and paths of such networks.

Distance function for links:

$d_L(i,i') \equiv d(L(i),L(i')) \equiv ?_j\, a_{i,j} + ?_j\, a_{i',j} - 2?_j\, a_{i,j}\, a_{i',j}$, i and i' in {1,…,I}.

Distance function for paths:

$d_T(j,j') \equiv d(T(j),T(j')) \equiv ?_i\, a_{i,j} + ?_i\, a_{i,j'} - 2?_i\, a_{i,j}\, a_{i,j'}$, j and j' in {1,…,J}.

Lemma: The distance functions above satisfy non-negativity, symmetry, and the triangle inequality.

Proof: $d_L(i,i')$ is the L_1 distance (as well as the Hemming distance and the squared Euclidean distance) between row i and row i' of the matrix A. Similarly $d_T(j,j')$ is the L_1 distance (and the Hemming distance and the squared Euclidean distance) between column j and column j' of A.

Comments, interpretations, and properties:

(i) Denote: W(i)={all paths passing through L(i)}, W(i,i') =P W(i)nW(i'), and let w(i) and w(i,i'), respectively, be the cardinality of these sets. Then:

$$d_L(i,i') = w(i) + w(i') - 2w(i,i') = w(i,i) + w(i',i') - 2w(i,i').$$

Similarly, denote: M(j)={all links in the path T(j)}, M(j,j')= M(j)nM(j'), and let m(j) and m(j,j'), respectively, be the cardinality of these sets. Then

$$d_T(j,j') = m(j) + m(j') - 2m(j,j') = m(j,j) + m(j',j') - 2m(j,j').$$

(ii) Note: $?_i\, w(i) = ?_j\, m(j) = ?_{ij}\, a_{i,j}$.

Manuscript received April 7, 2003. This work was supported in part by grants from the NSF, NSA, and DARPA.

Yehuda Vardi is with the Department of Statistics, Rutgers University, Piscataway, NJ 08854 USA (1-732-445-2692; fax: 1-732-445-3428; e-mail: vardi@stat.rutgers.edu) and with AVAYA Labs, Basking Ridge, NJ, USA.

(iii) Imagine that each link has a dedicated 'lane' for each SD path passing through it. Then w(i) is the number of lanes in L(i), and can be thought of as the "*width*" ("band-width", "capacity", etc) of L(i). Since w(i,i') is a count of the lanes that are common to both link L(i) and L(i'), it can be thought of as their *common width*. With similar interpretation, think of m(j,j') as the *common length* of T(j) *and* T(j'). Using this terminology we note that the distance between two links/paths is the sum of their norms (width/length) minus twice their *common* norm (width/length), respectively. When desirable, it is simple to rescale the metrics, non-linearly, to uniformly bound all distances, for instance, by the norm of the longest/widest path/link, respectively. See 'lemma' on rescaling below.

(iv) When two links/paths are close in distance, they likely share many common paths/links, respectively, and therefore will likely fail or operate flawlessly together. For instance, $d_L(i,i')=0$ holds if, and only if, the traffic-flows on L(i) and L(i') are identical. Thus, the proposed metrics are akin to measures of pair-wise statistical dependency on the network's components. To further demonstrate this property, note that when the traffic between all SD pairs are assumed to be independent Poisson with a common mean of unity, one gets: Variance(traffic on L(i) - traffic on L(i')) = Var(traffic on L(i)) + Var(traffic on L(i')) - 2Cov(traffic on L(i) , traffic on L(i')) = w(i) + w(i') - 2w(i,i') = $d_L(i,i')$.

Measuring distance between topologically similar networks: It is natural to compare traffic flows on topologically similar networks by comparing their routing matrices. Toward this end, for two networks with the same number of nodes (and hence pairs of nodes) and links, but possibly different routing matrices, say A and A' respectively, let **B**={all IxJ matrices B derived from A' by row- and column- permutations}. Define:

$d_N(A,A') \equiv \min\{\|A-B\|_1;\ B \text{ in } \mathbf{B}\}$ = L_1 distance between A (strung out as a vector) and the set **B**.

Lemma: argmin{||A-B||$_1$; B in **B**} = argmax{ ? $_{ij}$ $a_{i,j}$ $b_{i,j}$; B in **B**}.

Proof: For B in **B**, {||A-B||$_1$=? $_{ij}$ $a_{i,j}$ + ? $_{ij}$ $a'_{i,j}$ - 2? $_{ij}$ $a_{i,j}$ $b_{i,j}$.

Note: $d_N(A,A')$ is invariant under rows and columns permutations, so when two networks differ only in labels of links or nodes but otherwise have the same flow pattern, the distance between them is zero.

Rescaling: It might be desirable in certain applications to 'rescale' a distance function so it lies in a predetermined finite interval. This could be useful, for instance, if one is to bring into a common denominator the comparisons of networks of significantly different dimensions across different applications. To do so we note the following:

Lemma (Exercise 3.29 in Apostol (1974)): For any metric d(x,y), D(x,y) = d(x,y)/(1+d(x,y)) is also a metric (on the same space). Clearly 0=D(x,y)<1.

For completeness, we outline a proof: Symmetry, non-negativity, and equality to zero if and only if x=y, follow from the corresponding properties for d(,). To verify that D(,) satisfies the triangle inequality, let a=d(x,y), b=d(x,z) and c=d(z,y). Since d(,) satisfies the triangle inequality, we have a=b+c. Because of the monotonicity of d/(1+d) for d=0, we have: a/(1+a) = (b+c)/(1+b+c). Thus, it is enough to show that for any b,c=0, (b+c)/(1+b+c) = b/(1+b) + c/(1+c). Simple algebraic manipulations verify this.

Extensions and comments:

(a) *Topologically-different data-networks*: When two data-networks are topologically different, including different number of links or nodes, the corresponding routing matrices, say A and A', have different sizes. By far, this scenario would be the rule rather than the exception in practical situations. The most natural extension is to add the smallest number of rows and columns of zeros, as necessary, to A and/or A' to equate their size. Once they are embedded in the higher dimensional space their common size is [min(I,I')]x[min(J,J')]. This is like adding nodes or links that are disconnected from the network, and it is equivalent to embedding the routing matrices in the smallest higher dimensional space that contains both. We then apply the metric d_N(,) above to the (possibly) augmented matrices.

(b) *A measure of similarity for general networks*: Consider two networks, say with N and N' nodes and I and I' directed edges, respectively. We do not consider any notion of traffic flow, but rather think of the networks as static objects, i.e. a directed graphs. The question is: Can we extend the above concepts to a measure of similarity between the two networks? Here is a natural extension: For simplicity, assume first that I=I', N=N', and that the networks are strongly connected. In each network, for any ordered (SD) pair of nodes, say $j=(j_1,j_2)$, let T(j) be the *shortest* path from the "source" j_1 to the "destination" j_2. Assume, for now, that there is always a unique shortest path between any such SD pair (so the graph is strongly connected). This associates with each network, a unique "routing matrix" by assuming that traffic always flows on the shortest path between any SD pair. The two networks can now be compared using the metric $d_N(A,A')$ as defined above. While this approach is most reasonable for similar (in size and connectivity) networks, it can formally be extended to general graphs: To remove the assumption of equal dimensions, apply the embedding procedure as noted in (a) above. To remove the assumption of a unique shortest path between any SD pair, include a routing column for *each* of the shortest paths. To remove the assumption of connectivity between any pair of nodes, add columns of zeros for SD pairs that are not connected. To turn this into a similarity index, or rather a bounded metric, apply the 'rescaling lemma' above.

(c) *Challenges and practical considerations*: For large networks, there would be computational challenges in carrying out some of the 'modules' above, such as for instance, calculating the shortest path for *all* SD pairs. A practical approach could be to replace the "shortest path" with a "good path" using an efficient algorithm, repeat the process, and average across simulation experiments. Another challenge would be the minimization of $||A-B||_1$ (alternatively, maximization of $?_{ij}\ a_{i,j}\ b_{i,j}$) over *all* matrices B obtained from A' by permutations of rows and of columns. Perhaps an efficient 'proxy procedure' which maybe sub-optimal but nonetheless 'good', could be developed here.

(d) *Data-networks with Random routing*: Since random routing can be expressed as a probabilistic mixture of fixed routings one can measure distances in the random-routing case as a weighted average of the corresponding fixed-routing distance functions. (Note that a mixture of metrics is a metric.) This approach may have some theoretical value, but is computationally unrealistic. A more practical approach is to replace the zero/one routing-matrix A with a *frequency of utilization matrix* P: $p_{i,j}$ = P{L(i) in T(j)}= the probability that link i is included in the random path T(j). We then proceed to measure distances between the networks' links and, separately, its SD pairs, as the distances between the corresponding row pairs and, separately, column pairs, of $P=\{p_{i,j};\ i=1,\ldots I,\ j=,1\ldots,J\}$. Any reasonable distance function, such as Euclidean, L_1, and others could be used, and the choice should depend on the application. We note that for communication networks, there are hardware and software tools that can collect data and empirically estimate the $p_{i,j}$'s. (Note also the discussion on Markovian routing schemes in Vardi (1996), where the $p_{i,j}$'s are conditional probabilities, and not as defined above.)

3. Applications and Motivating Examples

In NISS (2003) Mark Crovella (see Lakinha et al, 2003) demonstrated that standard tools used in the exploration of network topology lead to severely biased estimation of the distribution of node degree. The nature of the bias appears to be such that the further away a node is from the launching cite of the trace-route probe, the less likely it is to be included in the sample. This is similar to size-biased sampling in statistics, but the lack of a formal distance function suitable for a communication network makes it impossible, or at least hard, to suggest a bias-correction method such as, for instance, in Vardi (1982, 1985). The distance functions proposed here could facilitate such a correction. The second application, presented by Bin Yu, also at NISS (2003), is that of maximum pseudo-likelihood estimation for internet tomography. In Liang and Yu (2003) a pseudo likelihood function is constructed for a fixed routing internet tomography problem. In their construction, the authors consider all possible pairs of rows of the routing matrix A, disregarding their dependencies, and use it to generate a "pseudo-likelihood" function from which a maximum pseudo-likelihood estimate (MPLE) is derived. As noted in our discussion above, two links (corresponding to two rows of A) may be close together (highly statistically-dependent) or far apart (highly independent). Accordingly, we can improve estimation by assigning weights so that each pair of rows gets a weight proportional to the distance between the two corresponding links. Such a weight-modified MPLE should be a more efficient estimator.

4. References

Apostol, T. M., (1974), **Mathematical Analysis: A Modern Approach to Advanced Calculus** (2nd Edition), Addison-Wesley Pub Co.

Anukool Lakhina, John Byers, Mark Crovella, and Peng Xie (2003), Sampling Biases in IP Topology Measurements, **Proceedings of IEEE Infocom**, San Francisco, California, April 2003.

Mark Coates, Alfred Hero, Robert Nowak, and Bin Yu (2002), Internet Tomography. **Signal Processing Magazine**, vol. 19, No. 3 (May issue), 47-65.

G. Liang and B. Yu (2003), Maximum Pseudo Likelihood Estimation in Network Tomography, **IEEE Trans. on Signal Processing** (Special Issue on Data Networks, to appear).

NISS (2003), Workshop on Internet Tomography, **NISS** (National Institute of Statistical Sciences) headquarter in NC, March 28, 2003.

Vardi, Y. (1982), Nonparametric estimation in the presence of length bias. **Annals. Statist. #**10, 616-620.

Vardi, Y. (1985), Empirical distributions in selection bias models. **Annals. Statist. #**13, 178-203.

Vardi, Y. (1996), Network tomography: Estimating source-destination traffic intensities from link data. **J. Amer. Statist. Assoc.**, vol. 91, pp. 365-377.