On the Internet Delay Space Dimensionality

Bruno Abrahao Robert Kleinberg*
Department of Computer Science
Cornell University
Ithaca, NY, 14850, USA
{abrahao,rdk}@cs.cornell.edu

ABSTRACT

We investigate the dimensionality properties of the Internet delay space, i.e., the matrix of measured round-trip latencies between Internet hosts. Previous work on network coordinates has indicated that this matrix can be embedded, with reasonably low distortion, into a 4- to 9-dimensional Euclidean space. The application of Principal Component Analysis (PCA) reveals the same dimensionality values. Our work addresses the question: to what extent is the dimensionality an *intrinsic* property of the delay space, defined without reference to a host metric such as Euclidean space? Is the intrinsic dimensionality of the Internet delay space approximately equal to the dimension determined using embedding techniques or PCA? If not, what explains the discrepancy? What properties of the network contribute to its overall dimensionality? Using datasets obtained via the King [14] method, we study different measures of dimensionality to establish the following conclusions. First, based on its power-law behavior, the structure of the delay space can be better characterized by *fractal* measures. Second, the intrinsic dimension is significantly *smaller* than the value predicted by the previous studies; in fact by our measures it is less than 2. Third, we demonstrate a particular way in which the AS topology is reflected in the delay space; subnetworks composed of hosts which share an upstream Tier-1 autonomous system in common possess lower dimensionality than the combined delay space. Finally, we observe that fractal measures, due to their sensitivity to non-linear structures, display higher precision for measuring the influence of subtle features of the delay space geometry.

Categories and Subject Descriptors: C.2.4 [Computer-Communication Networks]: Distributed Systems

General Terms: Measurement.

Keywords: Delay Space, Dimensionality, Network Embedding, Internet Structure.

*Both authors are supported by NSF Grant CCF-0643934.

1. INTRODUCTION

Network latency plays a central role in the design of a large class of Internet services as their performance is sensitive to the choice of the communicating participants. In light of this, coordinate-based network positioning systems have received considerable attention in the past few years [10, 12, 28, 32, 35, 39]. These approaches aim at providing a compact representation of the Internet delay space (i.e., the matrix of measured round-trip latencies between Internet hosts) by modeling the network as contained in a vector space. In this process, known as *network embedding*, each node is assigned a coordinate in a host metric space (e.g., Euclidean space) in such a way that the geometric distance between any two nodes estimates the real latency between them within a tolerable degree of error.

However, coordinate-based systems inherently suffer from embedding distortion, instability, slow convergence, and disappointing accuracy, as pointed out by [22] and [24]. Moreover, as discussed in [25] and [37], some aspects of the Internet graph make it difficult to model as a well-defined geometric object. As a result, these obstacles motivate positioning systems without coordinates [25, 26, 33, 44] as a more functionally viable alternative due to their improved accuracy, despite the fact that they are often measurement intensive and, in some cases, some types of queries are not available to network participants (e.g., the prediction of distances between any two other arbitrary nodes.).

As a way of understanding the potentialities and limitations of coordinate-based systems, we investigate a critical aspect influencing the effectiveness of this approach, namely the dimensionality of the Internet delay space. The main component of such systems consists of an embedding algorithm for which the number of dimensions of the host metric space (denoted hereafter by d) is a tunable parameter.

Embeddings with different numbers of dimensions result in different degrees of accuracy, since distance matrices posses a minimum intrinsic dimensionality [12, 28, 39]. In addition, since embedding techniques are based on some variation of an optimization problem aimed at minimizing the prediction error, the algorithms suffer from the *curse of dimensional-*

[1]A brief announcement to this work appears in the *Proceedings of the Twenty-Seventh Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing*, August 18-21, 2008, Toronto, ON, Canada.

[2]With the permission of the authors of [44], the processed 2385 × 2385 King dataset annotated with IP addresses, together with the datasets derived from it, is available at the Inetdim project website http://www.cs.cornell.edu/∼rdk/inetdim

ity. That is to say that the algorithm's complexity increases with the number of dimensions to the point that it becomes unable to deal with the overwhelming number of degrees of freedom to explore. Often times, this phenomenon occurs due to the unnecessary inflation of the metric space [6]. Finally, the convergence time, which increases with d, affects the stability of coordinates and the adaptability to changes, thereby affecting the reliability of the predictions [12, 22].

Previous studies indicated that embeddings of the Internet delay space can be created using 4 to 9 dimensions with reasonably low distortion in a Euclidean space [28, 39]. Dabek *et al.* [12] demonstrated that augmenting the embedding with *height* vectors, which are thought of as distance penalties incurred by traversing the last-mile access links in the Internet topology, allows one to use a lower-dimensional Euclidean space while retaining a similar level of accuracy. The application of Principal Component Analysis (PCA) [38] reveals the same dimensionality values [39].

However, this work prompts many questions which are still elusive: Can the dimensionality of the network be determined by embedding it into a host metric such as Euclidean space, or is there a more robust way of defining dimensionality, independent of the choice of host metric? Is the observed number of dimensions that produces low distortion embeddings in [12, 28, 39] optimal? More importantly, what properties of the Internet contribute to its dimensionality?

The characterization of the Internet delay space dimensionality, apart from its implications to the performance of coordinate systems, is by itself a topic of practical interest as it uncovers properties and opens new questions on the nature and complexity of the network [31].

As illustrated by [40], certain features of the Internet can be deduced purely from its delay-space geometry. For example, the partition of Internet hosts according to continents can be approximately reconstructed by clustering unlabeled distance data. As another example, it is unknown how much empty space is left by embedding the Internet using the current algorithms. If the empty space is significantly large, then extra overhead and complexity are being unnecessarily incurred by positioning systems. In addition, previous studies — as well as ours — have focused on datasets with up to a few thousand data points. However, theoretical results assert that in the worst case, the number of dimensions and distortion of an embedding increase logarithmically with the cardinality of the point set [8]. Therefore, a practical characterization issue, also critical for synthetic delay space generation purposes [45], is whether or not the dimensionality behavior and embedding distortion observed in previous studies will remain invariant with scaling to millions of nodes. Finally, the performance of positioning systems without coordinates also benefits from the characterization of the delay space geometry. For instance, the scaling guarantees of Meridian [44] are based on the assumption of a doubling metric whose main parameter is its dimensionality.

Tang and Crovella [40], and Huffaker *et al.* [16] demonstrated that geographic location (henceforth, *geolocation*) is a strong component to the Internet delay space. However, due to routing inefficiencies, caused by sub-optimal behavior of protocols, wide-area routing policies, and triangle inequality violations [39, 42, 45], great circle distances are not able to fully explain the Internet delay space dimensionality behavior. If this is the case, what other forces play a role in the dimensionality behavior?

This paper presents measurements and analysis to shed insight on these questions. We study the Internet dimensionality, defined as an intrinsic property of the distance matrix. This constitutes a geometrical invariant which does not refer to an external host metric space, such as Euclidean space, and does not include any structural distortions and unnecessary dimensionality inflation incurred by the embedding algorithms. We also study tools for exploring how and to what extent network properties drive the delay space dimensionality behavior. Using datasets obtained via the King [14] method, we compare four intrinsically-defined measures of dimensionality with the dimension obtained using network embedding techniques (such as Vivaldi [12]). We present three main conclusions. First, based on its power-law behavior, the structure of the delay space is best described by fractal measures of dimension, which measure a dimensionality intrinsic to the dataset, rather than by integer-valued parameters, such as the embedding dimension or PCA. Second, the intrinsic dimension is *much smaller* than the dimension predicted by the latter methods: in fact, by some measures, the intrinsic dimension is less than 2. Third, we quantify to what extent geolocation drives the dimensionality behavior and we present observations that suggest how the AS topology is reflected in the delay space. More specifically, we show that subsets of the data which can reach each other without going through a transit link between two Tier-1 providers consistently exhibit lower fractal dimension than the combined delay space, and that no such dimensionality reduction is achieved when partitioning according to geolocation. Given the properties observed above, we finally show evidence that fractal measures of dimensionality, due to their sensitivity to non-linear structures, display higher precision for measuring the influence of subtle features of the delay space geometry which are not captured by other dimensionality measures.

The rest of this paper is organized as follows. Section 2 presents the related work, Section 3 describes our experimental methodology, and Section 4 describes the notions of dimensionality we use in this work. Finally, Section 5 applies the proposed measures to study delay space features, and Section 6 offers our conclusions.

2. RELATED WORK

The first work to propose network embedding for the Internet was GNP (Global Network Positioning) [28]. In this work, Ng and Zhang tackle the complexity of the embedding by incrementally determining the coordinates of participants with respect to a set of a few previously chosen nodes (beacons or landmarks). These infrastructure nodes measure their inter-distance and determine their coordinates in some d-dimensional metric space. As a result they function as references in space so that subsequent incoming nodes can determine their own coordinates. This is done via a non-linear optimization problem that aims at minimizing the overall discrepancy between the geometric and measured distances. Hence, arriving participants compute the same minimization problem to determine their own absolute coordinates, by actively measuring their distance to the already-oriented beacons. Later on, a theoretical justification for the success of this approach was given by Kleinberg, Slivkins, and Wexler [19]; we discuss their work in greater detail below. Other examples of systems with similar scope are [10], [32], and [35].

On comparing the distributions of relative errors incurred by GNP using different number of dimensions, Ng and Zhang indicate that for the dataset studied, the best results were achieved using 7 to 9 dimensions.

Tang and Crovella [39] address the complexity and cost of network embedding by proposing the Lipschitz embedding. This method relies on the assumption that, by the triangle inequality, two nearby points, say a and b, have similar distance to a third point x, that is $|d(a,x) - d(b,x)| \leq |d(a,b)|$ and is defined in terms of a set D of subsets of a point set X. The distance function d defines the distance from a point x to one of the sets $L_i \in D$ as the distance from x the nearest point of L_i. Thus, the embedding is the mapping $\phi(x) = [d(x, L_1), d(x, L_2), ..., d(x, L_{|D|})]$, where $|D|$ corresponds to the dimensionality of the embedding. When every L_i is a singleton, each element represents a beacon and, therefore, component j of vector $\vec{x_i}$ is actually the measured distance from x to landmark j.

Since the number of dimensions required by the above embedding is high, the authors apply PCA, discussed in Section 4.3.2, to determine an r-dimensional space (where $r \ll n$) in which the data can be approximated with low loss of accuracy. As a result, each $\vec{x_i}$ is transformed into a $\vec{y_i}$, where each component of the latter is a linear combination of the distances to landmarks. Thus, they can be seen as distances to virtual landmarks. The method is evaluated using several real Internet datasets, and the number of dimensions that capture most of the overall variation of the data is between 7 and 9, incurring mean relative errors of 8 to 25 percent, consistent across the different datasets.

Subsequently, Dabek *et al.* proposed Vivaldi [12], which is a coordinate-based system that uses a mass-spring relaxation problem to determine the coordinates of nodes. In spite of using beacons, like GNP, Vivaldi has the advantage of not requiring fixed nodes serving this role and provides a degree of accuracy that is competitive with that of GNP. More interestingly, the authors propose the idea of unidirectional height vectors to augment the geometric model. Intuitively, the core of the network is mapped into a vector space as before whereas the borders of the network are assigned heights that penalize the distances of access link traversals. As a result, nodes can be placed up or down in order to accommodate conflicting distances in low dimensionality. It has been shown that this approach was able to embed the dataset considered into 2-space plus heights with competitive accuracy as compared to the embedding into 5-space.

The Vivaldi project also explores alternative geometric spaces, such as spherical and cylindrical, given that curved spaces resembles the surface of the globe around which the Internet is deployed. However, since most of the core links are centered in the U.S. and Europe, and due to the fact that there is no communication passing through the poles, the Internet does not wrap around the earth and, therefore, it has been shown that these approaches are no better than simply fitting the network into a plane space. Shavitt and Tankel explore embeddings in hyperbolic space which accommodate distance conflicts in low dimensional spaces, achieving a similar level of success as the height vectors [36].

More recently, studies quantified the inaccuracy produced by current positioning systems. Ledlie, Gardner, and Seltzer demonstrated that the relative errors increase with the cardinality of the set of hosts, and that the convergence to and maintenance of stable coordinates produced by the embedding algorithms are barriers to the effectiveness of such systems [22]. Lua *et al.* confirm that the degree of inaccuracy is beyond tolerable and propose new methods to quantify this fact, which is otherwise hidden by analyzing cumulative distributions of relative errors [24].

Zhang *et al.* characterized the delay space and proposed a synthetic data generator that improves upon existing topology generators [45]. Later on, the same group studied the impact of triangle inequality violations found in the delay space on overlay networks [42].

There is a rich body of theoretical work on questions regarding the existence of low-distortion embeddings of finite point sets into Euclidean space and other host metrics, as well as the computational complexity of algorithms for computing such embeddings. The starting point for much of this research is Bourgain's famous theorem which asserts that every metric space of cardinality n may be embedded with distortion $O(\log n)$ in a Euclidean space of dimension $O(\log n)$ [8]. A randomized algorithm for computing such an embedding was supplied by Linial, London, and Rabinovich [23]; the algorithm is based on semidefinite programming combined with a random-projection method due to Johnson and Lindenstraus [18]. Although there has been progress on the problem of minimizing the *additive distortion*, i.e., the maximum additive error over all pairs of points, the corresponding problem of computing low *multiplicative distortion* embeddings into lower-dimensional spaces — e.g., Euclidean spaces of dimension $o(\log n)$ — is an algorithmic problem of daunting complexity. Indeed, it is NP-hard to approximate the minimum-distortion embedding of an n-point metric into a d-dimensional Euclidean space within an approximation factor less than $\Omega(n^{1/12})$ [17]. Finally, several recent papers have considered the problem of computing *embeddings with slack* [5, 19], in which an ε fraction of all distances may be arbitrarily distorted and the rest must satisfy a low-distortion guarantee. This work, which uses beacon-based embedding techniques *a la* GNP, culminated in a theorem that every finite metric space admits an embedding in $O\left(\log^2 \frac{1}{\varepsilon}\right)$ dimensions with $O\left(\log \frac{1}{\varepsilon}\right)$ distortion and ε slack. Note that both the distortion and the dimensionality of the host metric in this theorem are still too high to be of practical value for network coordinate systems.

In Section 4 we study the power-law behavior of the delay space and its relationship with fractal dimensions. Many different power-laws and self-similar phenomena have been documented in the literature on Internet measurement. Examples include power-laws in the distribution of packet rates on an Ethernet link [43], inter-arrival times for FTP connections and TELNET packets [30], HTTP connections [11], different aspects of the Internet topology [13], and round-trip measurements in a time series of pings between a single pair of hosts [4].

The measures of fractal dimensions used in this work were also used by Belussi and Faloutsos [7]. Their work demonstrated that when spatial datasets behave like fractals (defined in Section 4.2.1) over a wide range of distances, one can use measurements of their fractal dimensions to rapidly estimate the spatial selectivity in range queries.

To the best of our knowledge, our work is the first to study the underlying dimensionality of the Internet delay space as a separate issue from the embedding of this delay space in any particular metric space, the first to document power-laws in the Internet delay space and to apply fractal geometry to its characterization, and the first to explore the impact of the Internet's AS-level topology on its delay-space geometry. The dimensionality revealed by our techniques is significantly lower than the dimensionality of the embeddings used in the prior work [12, 28, 39], although the existence of an algorithm that produces embeddings with this low dimensionality behavior is still an open question.

3. METHODOLOGY

Our main analysis is based on the Meridian dataset presented by Wong, Slivkins, and Sirer [44], and some of our findings are also supported by observations made using the MIT King dataset [2, 12].

The Meridian dataset was collected between May 5-13 2004 via the King [14] method, containing latency measurements between more than 5200 DNS servers. The list of sites to measure was determined by randomly picking website names from a set of 593160 entries obtained from the DMOZ and Yahoo directories. The raw data consists of a set of asymmetric measurements between pairs of DNS servers, that is, the RTT's in microseconds between two servers A and B, measured by recursively querying A for domains served by B, and vice-versa. The number of asymmetric measurements per pair varies between 1 and 20 entries with median 11. In order to create the matrix used in this work, we took the union of the asymmetric measurements for each pair, thereby making the dataset symmetric. We subsequently filtered out the pairs with less than 10 measurements, in order to minimize biases due to queuing delays at routers or DNS servers, and then computed the median of the symmetric measurements for the remaining pairs. Finally, we approximate the largest clique in the resulting incomplete matrix via a 2-approximation algorithm for the vertex cover problem (i.e., to eliminate the missing entries by removing the minimum number of nodes), resulting in a all-pairs matrix with 2385 hosts, annotated with their IP addresses (henceforth referred to as "IPs" for brevity).

The MIT King dataset was first used to study Vivaldi's behavior [12]. It was also collected using the King method and contains measurements among 1953 hosts, selected by finding the NS records of IP addresses of participants in a Gnutella network. Nevertheless, after applying the above data cleaning process, only 298 nodes remained in the dataset. Although the size of this dataset does not allow us to use it for analyzing all the aspects discussed in this work, it can still be used to support some of our findings.

We also merged the delay data in the Meridian dataset with the underlying AS topology by obtaining a snapshot, from the same period the delay dataset was collected, of the customer-provider AS graph from the CAIDA AS relationships dataset [3]. Using the combined data, we made a decomposition of the network into AS trees, each rooted at one major Tier-1 Autonomous Systems (AS). Accordingly, each piece comprises the Tier-1 network itself, together with its downstream network of AS customers. After decomposing the whole network in this fashion, we classified the IPs found in the Meridian dataset into each piece. The number of IPs found in each piece is summarized on Table 1.

Notice that the sum of the number of IPs in each network exceeds the total number of IPs in the dataset. This is because most of the IPs are located in multihomed networks (i.e., are served by multiple providers). In fact, according to our data, more than 60% of the customers have contracts with more than one provider, and the number of upstream providers per network can be as high as 13. Thus, our decomposition does not consist of a partition of the space, and, in fact, some of the subsets contain more than 50% of the nodes in the whole clique. [3]

Table 1: List of the major Tier-1 AS together with the number of IPs in their downstream networks represented in the Meridian dataset.

AS		Meridian
AS#	**Name**	**Hosts**
2914	NTT Comm.	1212
209	Qwest	1227
3561	SAVVIS	1389
3356	Level 3	1454
7018	AT&T	1487
3549	Global Crossing	1515
701	Verizon	1529
1239	Sprint Nextel	1604

We emphasize that the King method is a convenient way to obtain vantage points using DNS servers, which are generally well-connected hosts. Thus, although its hosts are geographically and topologically diverse, it can be argued that datasets collected via King could only give us an approximate picture of the delay space geometry as composed of core and edge networks. Furthermore, both datasets contains violations of the triangle inequality to a degree consistent with that found in the characterizations by [12], [39], [42], [44], and [45].

The unavailability of Internet latency datasets is a major limitation to the study of the dimensionality of the delay space. We have investigated the other publicly available datasets of this kind but they are either limited in scale, i.e., do not contain a large enough all-pairs matrix so that our analysis can be applied, or they are not annotated with IP addresses, which makes the above decomposition impossible.

In order to analyze the geographic component of the delay space, we queried the *hostip.info* database [1] for the IPs contained in the Meridian dataset, obtaining their latitude and longitude at the time of writing. Since these IPs belong to DNS servers of large domains that are not as likely to have their IPs reassigned as are smaller domains, we resort to the assumption (not quantified) that a large fraction of the IPs contained in the Meridian dataset were not reassigned since 2004. Even though we believe that this is a reasonable approximation, it should be emphasized that, together with the fact that geolocation is currently a process with high degree of inaccuracy, this could combine several sources of error for the particular set of results in Section 5.1.

[3] A recent result indicates that a single snapshot of the inferred AS topology map is believed to miss around 10% of the customer-provider links involving Tier-1 and Tier-2 networks [29]. While this does not lead to missclassification, it could possibly exclude the membership of some domains (and its downstream customers) in some of the pieces.

Figure 1 presents a coarse-grained visualization of the geolocation of nodes in the Meridian dataset.

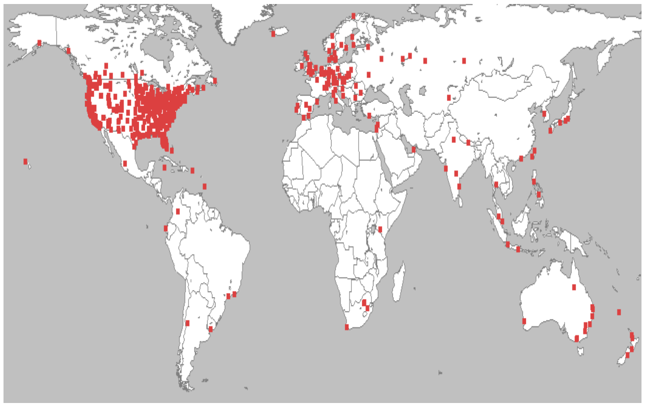

Figure 1: Geographic location of nodes in the Meridian dataset.

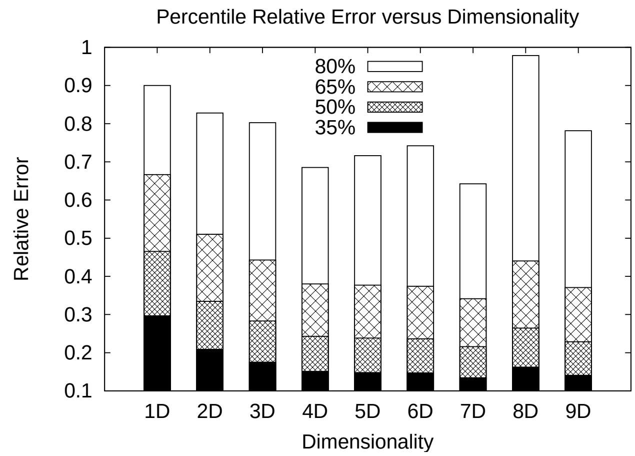

Figure 2: Percentiles of relative errors produced by Vivaldi using different values of d.

4. DIMENSIONALITY MEASURES

As a starting point for introducing the measures of dimensionality that we use in this work, let us consider the following problem. Suppose that a surveyor chooses a set X of 2500 random points in the plane and measures the distances between all pairs using a method that introduces 5% relative error due to measurement noise. Given the matrix of measurements, but not the coordinates of the actual points, how could one deduce that the data came from a point set in 2 dimensions, rather than 1 or 3? We consider this problem further in the following sections.

4.1 Embedding dimension

An obvious answer to the question question posed in the previous section is: for $d = 1, 2, 3, \ldots$, try to embed the points in d dimensions using an embedding algorithm such as Vivaldi. Stop at the lowest dimension D which permits an embedding with small quartiles of relative errors and let D denote the *embedding dimension* of the dataset.

We applied this process to the Meridian dataset by embedding the network into Euclidean space using Vivaldi, available as part of the P2Psim package [2]. We varied the number of dimensions from $d = 1$ to $d = 9$ and Figure 2 presents the outcomes of this experiment by displaying the 35-th, 50-th, 65-th, and 80-th percentiles of relative errors incurred by embedding the whole delay space with different values of d (with the percentile values chosen in such a way that the discrepancy of the distributions could be well captured). In this plot, we can observe that there is a fast improvement in accuracy up to $d = 4$ and a slow improvement up to $d = 7$. Surprisingly, after 7 dimensions, the accuracy of the algorithm gets worse, exposing the threshold beyond which the curse of dimensionality starts to affect the algorithm's performance.

A benefit of this approach is that, if successful, it actually recovers the coordinates of the original points (up to translations and rotations).

However, it also has many drawbacks: first, embedding algorithms are slow, even for relatively small values of d. Second, finding an embedding that minimizes distortion is computationally intractable in the worst case [9]. Last but not least, if the measured distances reflect a metric other than the Euclidean distance (e.g., the hyperbolic metric or the Manhattan metric), the algorithm may fail to find a low-distortion embedding in any dimension.

Moreover, by attempting to fit the distances precisely, the algorithm may produce a high-dimensional embedding with lots of unnecessary empty space, obscuring the fact that the points of the embedding really lie in a lower dimensional subset of that space. For example, if a point set X were located on a hilly terrain instead of a flat plane, the algorithm would output an embedding using 3 dimensions despite the fact that all of the data lies along a 2-dimensional surface in 3-space.

Finally, the embedding may fail to reveal lower dimensional substructures which constitute important features of the distance matrix. For example, suppose that the entries of the distance matrix are estimates of the time required to walk between various locations in an office building with several floors. The geometry of the office building can be most accurately modeled as a small number of 2-dimensional pieces (the floors) with a small number of "gateways" (the stairwells) connecting these pieces together. Embedding the distance matrix accurately in a Euclidean space would require at least 3 dimensions — probably more, since shortest paths in the office building are very different from shortest paths in 3-space — obscuring the inherent 2-dimensionality of the office building's floor plan. Like the office building, the Internet delay space is also composed of smaller pieces (autonomous systems) which meet only at prescribed gateways (customer-provider and peering links). When representing the geometry of the Internet delay space, one should not choose a representation which obscures this intricate structure.

For purposes of estimating the dimensionality of a point set there are several other, more lightweight, ways of defining dimensionality using structural properties of the distance matrix itself. These estimates can be done without making reference to an outside "host metric", such as Euclidean space, and without computing coordinates to represent its points.

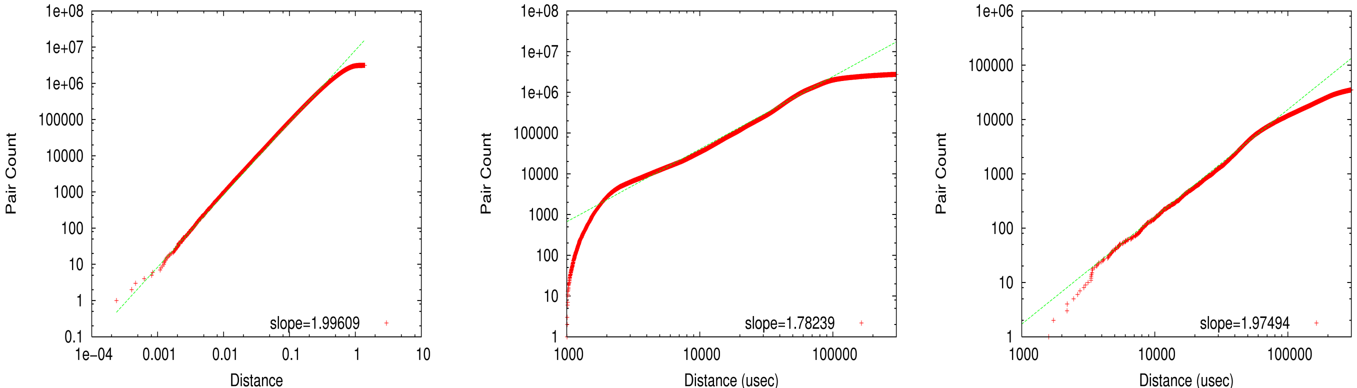

(a) Unit Square, $D_2 \approx 2$ (b) Meridian matrix, $D_2 = 1.78$ (c) MIT King Matrix, $D_2 = 1.97$

Figure 3: Correlation fractal dimension, D_2, of a) a set of 2500 random points in a unit square, and the Internet delay space as represented by b) the Meridian matrix and c) the MIT King matrix.

These methods also capture the effects of the intricate patterns mentioned above. We next introduce these definitions of dimensionality, explaining the applicability of each and their degree of accuracy for the solution to the problem introduced in the beginning of this section. Although the following notions of dimensionality were introduced in the theory of metric spaces (which assumes the triangle inequality), all of them have the desirable property that they are applicable even in datasets (like ours) which contain triangle inequality violations, often yielding meaningful results.

4.2 Correlation Dimension

Suppose we pick a random point $x \in X$ and a radius r, and we count the number of other points whose distance from x is at most r. If the points are random samples from a bounded region in the plane, then the expected number of such points in a region is proportional to its area. Hence the number of points within distance r of x should be proportional to r^2. For the same reason, in d dimensions, the number of points would be proportional to r^d. Hence, upon plotting the radius r against the number of pairs whose distance is at most r in logscale (which we denote as the *pair-count* plot), for special point sets which produce a straight line over a given range of interest (i.e., exhibit power-law behavior), we can interpret the exponent of the power-law d (i.e., the slope of the line) as reflecting the underlying d-dimensionality of the dataset. We refer to d as the *pair-count exponent* [7], which corresponds to the correlation fractal dimension, D_2 [34], further discussed in Section 4.2.1.

Figure 3(a) illustrates the pair-count plot of a set of 2500 random points in a unit square surface. Notice the presence of a power-law that persists over three decimal orders of magnitude. Notice further that the exponent of this power-law is almost exactly equal to the dimension of the space from which points were sampled, in accord with the theoretical prediction sketched above. In fact, for all "Euclidean objects", i.e., distance matrices obtained from uniformly-random point clouds in Euclidean d-space, the fractal dimension matches the Euclidean dimension.

Figures 3(b) and 3(c) present the pair-count plot of the Internet delay space as represented in the Meridian and MIT King datasets respectively.

The first striking feature of these plots is a power-law that persists roughly over two orders of magnitude, i.e., from 3ms to 100ms. (Note that this range of latencies includes almost every Internet route that is not trans-oceanic.) As a result, the Internet delay space exhibits the desirable property that it can be measured by fractal measures of dimensionality (see Section 4.2.1). The second unexpected observation is that the magnitude of its dimensionality is less than 2, represented by the pair-count exponents $D_2 = 1.782$ and $D_2 = 1.975$ in the Meridian and MIT King dataset respectively. The estimation of these values contains errors to a degree that would not affect the conclusions derived in this work. These dimensionality values are much smaller than the embedding dimension indicates (i.e., between 4 to 7 dimensions), suggesting a different geometric picture of the structure of the Internet delay space. Sections 4.3.2 and 5.2 discuss the reasons for this discrepancy. Finally, the power-law behavior, including the dimensionality value, is consistent across random subsets of the data, as discussed in Section 5.2.

The fractal measures can help us understand the weaknesses of embedding algorithms by showing how they affect the properties of the original delay space. Accordingly, upon computing the pair-count plot of the embedded network produced by Vivaldi in 7 dimensions, we can observe how the resulting coordinate space does not preserve the geometric properties of the original delay space and suffers a major dimensionality inflation. Figure 4(a) shows the resulting pair-count plot of the delay matrix reconstructed from the 7-space coordinates. Notice that the curve has a concave shape, thereby deviating from the power-law behavior of the original space. Moreover, the best effort to measure its dimensionality, by finding the best straight line fit to the curve, results in a pair-count exponent of value 5.46.

Although the Internet hosts live in a sphere that can be described by two coordinates in spherical space, the values near 2, which the fractal measures revealed for the delay space dimensionality, are not a reflection of the 2-dimensional structure of a sphere's surface. This phenomenon would be further explored in Section 5.1.

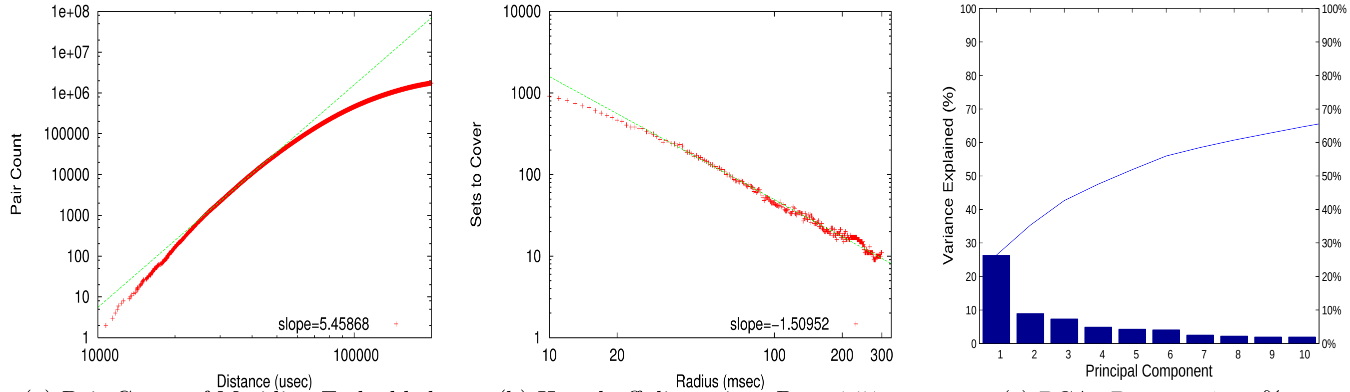

(a) Pair-Count of Meridian Embedded (b) Hausdorff dimension, $D_0 = 1.51$ (c) PCA, $D_{PCA} = 42.66\%$

Figure 4: Measures of dimensionality applied to the Meridian matrix: a) pair-count plot of the embedded Meridian data into 7-Euclidean space, b) the Hausdorff dimension computed by the greedy set cover plot, and c) Principal Component Analysis (PCA) applied to the Meridian matrix.

4.2.1 Fractal dimensions

The previous section introduced a measure of dimensionality which was based on measuring the exponent of a power-law arising in distance data. In general, this power-law does not necessarily arise and, when it does, it need not have an integer exponent. Point sets whose pair-count plots display a power-law are called *fractals*.

The correlation dimension is just an example from among an infinite family of fractal dimensions D_q, indexed by a non-negative number q. Formally, if μ is a fractal measure on a set Y and $A_1, A_2, \ldots, A_N$ is a partition of Y into pieces of diameter less than r, with $\mu(A_i) = p_i$ for $i = 1, 2, \ldots, N$, then[4]

$$D_q(Y) = \frac{1}{q-1} \lim_{r \to 0} \frac{\log\left(\sum_{i=1}^{N} p_i^q\right)}{\log r}. \tag{1}$$

The correlation dimension corresponds to the fractal dimension D_2. Among these dimensions, only the first three can be efficiently measured in practice.

Fractals may arise by applying recursive constructions in which a self-similar point set is composed of finitely many pieces, each of which is a scaled-down copy of the entire point set. If one samples a large number of points uniformly at random on such a perfect infinite fractal, and measures any of the fractal dimensions (e.g., for all q), they must coincide. For example, if a point set is made up of 2^p pieces, each of which is a copy of the entire set scaled down by 2^k, then all its fractal dimensions equal p/k. However, if the data is non-uniformly distributed inside the fractal, one gets a *fractal measure* or *multifractal* — a fractal together with a probability measure expressing the density of points at different locations.

Despite this recursive definition, the most surprising examples of fractal behavior are non-recursive structures commonly found in nature. Some examples are snowflakes, coastlines and the surface of the human brain [34]. The question of exactly what features of the Internet delay space lead to its fractal behavior is still elusive. In the search for these properties, we discovered some hints that are discussed in Section 5.2. However, this question is currently a subject of our ongoing work.

[4]The case $q = 1$ is exceptional. In equation (1) when $q = 1$, one uses the log of the entropy of the distribution $\{p_i\}$ in the numerator and drops the constant $1/(q-1)$ out front.

4.3 Other dimensionality measures

In this section, we introduce another instance of fractal dimension, namely *Hausdorff dimension* (D_0) and two dimensionality reduction techniques, namely *Principal Component Analysis (PCA)* and Isomap, explaining the relevance of each of them to this work.

4.3.1 Hausdorff dimension

Consider partitioning a point set X into low-diameter subsets. If X lies in a bounded region of the plane, then for every $r > 0$ it can be partitioned into $O(1/r^2)$ subsets of diameter less than $2r$, for example using grid cells of side length r. The analogous low-diameter covering in dimension d uses $O(1/r^d)$ subsets. Even if we are given only the distance matrix — so that it is infeasible to identify the partition into grid cells — a partition into low-diameter sets can still be constructed by considering the collection of all radius-r balls and selecting a sub-collection using the greedy set cover algorithm. For d-dimensional Euclidean objects the cardinality of this greedy covering will also be $O(1/r^d)$ with high probability, though it is less obvious than in the case of the grid-cell covering.

This suggests defining $N(r)$ to be the minimum size of a partition of X into pieces of diameter less than $2r$, and plotting r against $N(r)$ in logscale. For d-dimensional Euclidean objects we have seen that this will lead to a line of slope $-d$. For any distance matrix, if a power-law with exponent $-d$ is present over a given range of interest, we refer to d as the Hausdorff fractal dimension [34], or D_0, from the fractal definition presented in Section 4.2.1, Equation 1.

Figure 4(b) presents the measure of D_0 for the Internet delay space as represented in the Meridian matrix. Similarly to the pair-count plot, we observe the presence of a power-law with exponent -1.51.

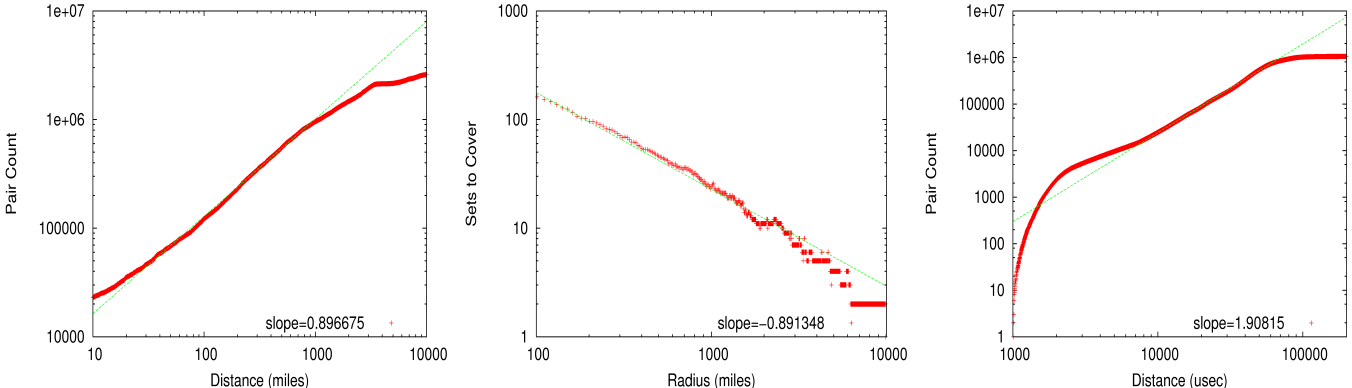

(a) Correlation dimension, $D_2 = 0.90$ (b) Hausdorff dimension, $D_0 = 0.89$ (c) Diameter-based Cluster

Figure 5: Dimensionality of the distance space based on great circle distances via a) the correlation dimension and b) The Hausdorff dimension. c) The correlation dimension of one of the latency-based clusters.

While D_2 encompasses the geometric structure of the delay space and also the spatial distribution of points, D_0 includes only the former notion [7, 34]. Nevertheless, we introduce D_0 for two main reasons. First, it exhibits the same power-law behavior displayed by D_2 when applied to the delay space as represented by the datasets considered (albeit with a different power-law exponent). This observation reinforces the fractal behavior of the delay space. Second, D_0 behaves similarly to D_2 when measuring networks that possess different intrinsic dimensionality values. Section 5.2 presents this behavior in greater detail.

Finally, a number of algorithms are known to run efficiently on metric spaces that possess some form of bounded *doubling dimension* [15]. In general, that is to say that every ball of radius R in the metric can be covered by at most 2^d balls of radius $R/2$. Similarly to the Hausdorff dimension, a d-dimensional metric space has doubling dimension approximately d. While the doubling dimension has convenient algorithmic properties, it is difficult to precisely determine the doubling dimension of a dataset. (See [21] for an example of this process.) The difficulty arises because, unlike the correlation or Hausdorff dimension which are statistically robust against perturbing a few entries of the distance matrix, the doubling dimension is very sensitive to outliers as it involves taking the maximum covering number over all balls in the entire metric space.

4.3.2 Principal component analysis

PCA is the main technique adopted in previous work to characterize the Internet delay space dimensionality [12, 22, 39]. In our experiments we have applied PCA directly to the distance matrix. This is similar to the analysis applied in [39], where the method was justified with the additional observation that the columns of the distance matrix itself represent the coordinates for a particular embedding of the delay matrix, namely the Lipschitz embedding in the L_∞ norm, discussed earlier in Section 2.

Figure 4(c) presents the application of this approach to the delay space as represented by the Meridian matrix. The common practice suggests that the dimensionality of a dataset is determined by the point d in the x-axis in which the contribution of the corresponding component differs significantly from the previous one, and the percentage variance explained by components with labels greater than d becomes negligible. However, as in the case of this plot, it is not always clear where to establish this threshold. Since we use PCA for the purpose of comparison of the dimensionality values of different networks, we define our PCA measure as the percentage variance explained by the three most significant components. Accordingly, low dimensionality implies more variance being captured by these components whereas high dimensionality tends to spread the variance across a greater number of components. In the case of Figure 4(c), the PCA measure equals 42.66%.

As PCA is a method grounded in linear algebra, when applied directly to a distance matrix, it is oblivious to non-linear relationships between different dimensions. [5] For example, if the points are sampled uniformly at random from a circle in the plane, PCA applied to the distance matrix will strongly indicate a 2-dimensional dataset despite the fact that all of the points belong to a 1-dimensional curve. Here is where the fractal measures come into play: by capturing non-linear, as well as linear, relationships among the dimensions, the fractal measures are able to capture patterns in the data otherwise ignored by PCA. Therefore, they provide a more genuine characterization of dimensionality.

4.3.3 Isomap

Finally, we apply Isomap [41], a geometric dimensionality reduction technique, proposed in the Machine Learning community, which is sensitive to both linear and non-linear correlations between the dimensions. Like PCA, when applied to a dataset, Isomap outputs the fraction of the total variance explained by each of the dimensions and produces a d-dimensional non-linear embedding where the d is a tunable parameter. Our results indicate that, similarly to the fractal measures, Isomap displays higher sensitivity to Internet structural properties as compared to PCA and the embedding dimension (see Section 5.2), thereby indicating that the delay space is rich in non-linearity. The Isomap results

[5] PCA can in principle capture non-linear relationships when combined with kernels [27]

presented in this paper were produced using the code that implements Isomap available on the authors' websites [41].

5. ON THE DELAY SPACE STRUCTURE

In an attempt to understand the features of the delay space that contribute to its dimensionality behavior, we study how and to which extent network properties, such as the geographic structure and some features of the AS topology, are reflected in its delay space dimensionality. The analysis in this section also demonstrates the applicability and effectiveness of each dimensionality measure in capturing these properties.

5.1 The Geographic Component

Our first experiment aims at quantifying the impact of the Internet's geographic structure on the delay space. For this purpose, we computed the great circle distances (in miles) for every pair of nodes and generated a new distance matrix. The first observation is that the pair-count plot, presented in Figure 5(a), exhibits a power-law that also persists for approximately two decimal orders of magnitude. As a consequence, it can also be measured using the correlation dimension and has exponent $D_2 = 0.897$. The plot for Hausdorff dimension, shown in Figure 5(b), is less conclusive, although it can also be approximated by a straight line with slope $D_0 = 0.891$. As expected, PCA applied on the distance matrix results in the two first components explaining 100% of the variation in the data.

Note that the geographic dimensionality value is less than 1 while the surface of a sphere has dimensionality 2. This difference can be ascribed to the large empty spaces (i.e., oceans) and the non-uniform geolocation of nodes (i.e., dominant clusters in North America, Europe and Asia).

Since the dimensionality of the geographic space is significantly smaller than that of the delay space, the contribution of geolocation does not fully explain the delay space dimensionality, albeit it is indeed a strong component, in accordance with the analysis by [16] and [39]. However, the fractal measures of dimensionality allow us to quantify the extent to which geolocation contributes to the overall delay space structure. Furthermore, it shows that the fractal behavior of the delay space is in fact present in the underlying geodesic space.

5.2 Dimensionality Reducing Decomposition

On searching for a (possibly recursive) structural feature of the delay space that could explain its fractal behavior, we discovered a dimensionality shift that can be ascribed to the structural configuration of the AS topology.[6]

Our intuition is based on the observation that peering points, especially those corresponding to transit links between two Tier-1 networks, are contained on a significant fraction of the Internet routes joining nodes which are downstream from different Tier-1 providers. Even though there is an abundance of peering points between non-Tier-1 networks, and multihomed domains are the norm, the significance of paths traversing transit links between two Tier-1

[6]In the interest of space, we present only a subset of the graphs from which we derived the conclusions in this section. However, the remaining graphs corresponding to all results presented here can be found in the companion website at http://www.cs.cornell.edu/∼rdk/inetdim

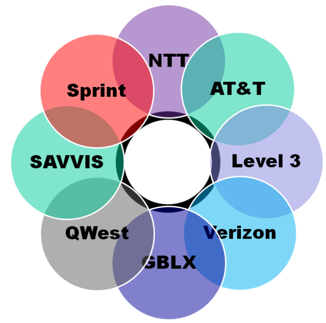

Figure 6: Network decomposition into intersecting pieces, each corresponding to a Tier-1 AS together with its downstream network.

networks could still have a major influence on the geometry of the Internet delay space. Thus, we ask the following question: what is the geometric effect of analyzing each Tier-1 AS downstream network in isolation, thus removing the distances that contain a contribution from the traversal of Tier-1 transit links?

We expected that this process would result in subnetworks whose delay space geometry is better behaved, as compared to their superposition, in which the subspaces corresponding to the different Tier-1 AS downstream networks would primarily connect to one another in the Tier-1 peering points, as illustrated in Figure 6.

An example of this scenario was previously hinted in [28] where the geometry exhibited by a more homogeneous network, when observed in isolation, is better behaved, as compared to a more diverse network. In the case of this study, the embedding of the *Abilene* network, connecting hosts through a fast Internet2 backbone, resulted in gains in accuracy as high as 40% in relative errors, using the same number of dimensions as in the embedding of a global network. This improvement was ascribed to the overall short distances of the paths.

In order to quantify the extent to which transit links are reflected in the Internet delay space, we measured each subnetwork (defined according to the decomposition presented in Section 3) using embedding dimension, correlation dimension, Hausdorff dimension, and PCA.

In order to confirm that the dimensionality reduction observed in this experiment is attributable to effects arising from the AS-level topology of the Internet, and not some other simpler explanation, we tested two alternative hypotheses: that the dimensionality of the delay space can be reduced by decomposing it into pieces of smaller *cardinality*, or that it can be reduced by decomposing into pieces of smaller *diameter*.

The first hypothesis is partially justified by Bourgain's theorem [8] that the Euclidean distortion of a metric space grows logarithmically with its cardinality, in the worst case. To test the hypothesis that cardinality reduction leads to dimension reduction, we created 121 random subsets of the whole matrix, each containing 1454 hosts (the median cardinality among all subnetworks). In Table 2, we denote by *random x%* the statistics derived from these random subsets corresponding to the x-th quartiles of these values.

Table 2: Dimensionality measures: correlation dimension (D_2), Hausdorff dimension (D_0) and PCA found for the different AS networks.

	Dimensionality		
Network	D_2	D_0	**PCA**
Meridian	1.783	1.510	42.66
1239	1.780	1.333	46.54
209	1.637	1.225	28.42
2914	1.618	1.161	44.25
3356	1.691	1.230	48.60
3549	1.634	1.265	48.96
3561	1.686	1.239	49.54
7018	1.710	1.304	49.68
701	1.701	1.244	49.68
random 0%	1.725	1.389	27.71
random 25%	1.762	1.465	45.76
random 50%	1.775	1.506	46.67
random 75%	1.789	1.554	53.10
random 100%	1.837	1.682	72.86

To test the second hypothesis, that diameter reduction leads to dimensionality reduction, we created a new decomposition of the distance matrix by the following process. Starting from an element selected from a geographically diverse set of hosts, we grew a ball around it by selecting its 1454 closest neighbors, and examined the delay space consisting of all inter-distances among these 1454 hosts. We have constructed 12 of these subsets centered at hosts located in 9 different countries, 4 continents.

Table 2 summarizes the dimensionality of each the subnetworks and values of the statistic thereof for the random subsets[7] in terms of the correlation dimension, Hausdorff dimension, and PCA measures. The corresponding measures for the whole network are also displayed in the table for reference, under the label *Meridian*.

The first observation is that the power-law behavior observed in the whole matrix was preserved over the same range of distances in the submatrices, though not necessarily with the same exponent. Second, with the exception of the correlation dimension of the subnetwork rooted at AS 1239 *(Sprint)*, all other networks exhibit *smaller* dimensionality than the whole matrix according to the two fractal measures.

The values of both the correlation and Hausdorff dimensions for the random networks, including the minimum of these values (random 0%), are consistently higher than those of every subnetwork (except for the pair count of network 1239), and a higher percent of these values are close to that of the whole matrix. As one increases the size of the original dataset and the number of representatives in the set of random networks, one expects even less deviation between the correlation and Hausdorff dimensions of the original dataset and the median values measured in the random subnetworks.

The decomposition based on clustering of growing balls resulted in the following observations. For all clusters, the pair-count plots present a power-law persisting over one decimal order of magnitude less than the whole matrix. For smaller values of r (in the range 3ms to 10ms) the number of pairs at distance r is significantly *higher* than the number of pairs predicted by the power-law approximation. This is perhaps not surprising, since the clusters were selected for their proximity to a single central host. Figure 5(c) shows one example of this outcome for one of the clusters. Moreover, the pair-count exponents (i.e., correlation dimensions) computed over the range in which the power-law persists indicate dimensionality consistently *greater* than that of the whole matrix. The deviation from a power-law is also observed in the range from 20ms to 30ms of set cover plots. In addition, the Hausdorff dimension measured over the range from 30 to 90ms is consistently greater than that of the whole matrix for all clusters.

Interestingly, delay space dimensionality reduction cannot be explained by a corresponding reduction in the geographic component. Upon applying the same decomposition analysis to the great-circle distances of the subnetworks, we observe no statistically significant difference in the power-law exponent of the subnetworks as compared to the pair count and set cover exponents found for the combined network.

It is important to emphasize that the AS relationship graph is complex and, therefore, other forms of decomposition could result in pieces that display the optimal dimensionality reduction. Nevertheless, the analysis presented here shows evidence, with statistical significance, that the presence of *Tier-1* transit links has a non-negligible effect on the Internet's delay space geometry.

Interestingly, both fractal measures are sensitive enough to capture this structural change, in the form of a reduction in fractal dimension when one restricts attention to a subset of the delay space consisting of a single Tier-1 provider and its downstream networks. Table 2 also contains values for the percentage variance explained by the first three components (i.e., the three greatest in magnitude) found by PCA. Notice that, as opposed to the fractal measures, there is no clear distinction in dimensionality behavior between the subnetworks and the random sets as reported by PCA. We have also computed the embedding dimension of each network in Euclidean space using Vivaldi and did not observe that decomposing the network into subnetworks had any effect on the embedding dimensionality. In fact, similarly to PCA, in some cases, the dimensionality of subnetworks is reported as being much greater than that of the entire network.

We hypothesize three possible explanations for the insensitivity of PCA and Vivaldi to the network decomposition, which was otherwise very well captured by the fractal measures. First, PCA and Vivaldi are oblivious to nonlinear relationships in delay space. Second, PCA and Vivaldi try to represent non-linear relationships using linear ones. Therefore, the process of accommodating these discrepancies causes the dimensionality to inflate. Finally, PCA and Vivaldi are suited to reporting an integer which summarizes the dataset's dimensionality; such methods are too coarse to detect a difference of 0.2 in the dimensionality.

To confirm our hypotheses we applied Isomap (see Section 4.3.3) to the networks. Figure 7 shows the outcome of this experiment. In the plot, the fraction of the total vari-

[7]The decomposition based on low-diameter subsets produced results which are incomparable with these results because the clustering rule significantly altered the nature of the power-law behavior, as explained further below. Hence the results of these experiments are not included in Table 2.

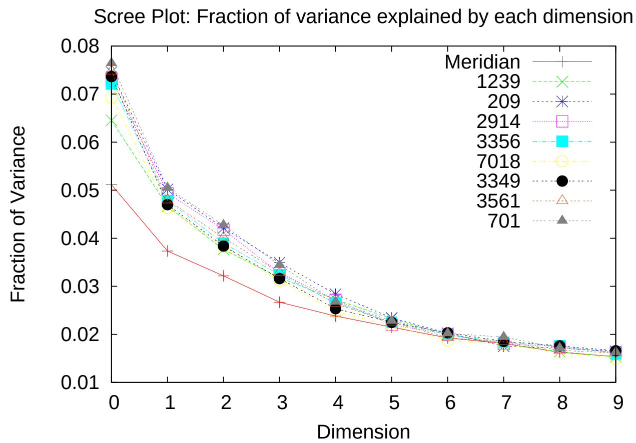

Figure 7: The scree plot outputted by Isomap.

ance explained by the first five (also five most significant) components of all subnetworks is consistently greater than that of the entire network. As opposed to what PCA and Vivaldi tell us, and in accordance with the fractal measures, the dimensionality of the subnetworks is indeed consistently *smaller* than the dimensionality of the whole network. However, the dimensionality values reported by Isomap for all networks that is still larger than those predicted by the fractal measures. This result shows evidence of non-negligible non-linear structures in the delay space and suggests that methods sensitive to them might reflect the structural properties of the Internet with superior degree of accuracy.

6. CONCLUSION

This work studies different dimensionality measures and investigates the dimensionality observed in the Internet delay space, thus providing insight into some of the underlying causes. Characterizing the geometry of the delay space is critical for designing and analyzing effective coordinate-based positioning systems and also sheds light on the structural properties of the network.

We have observed that the Internet delay space displays special structural and statistical characteristics in that it adheres to a power-law that extends over a significant range of Internet distances, namely the intra-continental distances. Therefore, its dimensionality can be characterized using fractal measures, sensitive to non-linear as well as linear structures in the delay space. Moreover, they are intrinsic properties of the delay space, independent of a target host metric space (e.g., Euclidean space). Therefore, the delay space dimensionality can be measured without computing coordinates for points, and is not subject to dimensionality inflation caused by embedding algorithms.

We have used the proposed fractal measures to quantify the extent to which geodesic distances are reflected in the dimensionality of the Internet, and we have also shown that, upon decomposing the Internet into subnetworks consisting of hosts that share an upstream Tier-1 autonomous system in common, we observe a reduction in the intrinsic dimensionality of the pieces.

Moreover, both fractal measures were able to capture the role of the Internet's AS-level topology in determining its delay space geometry, a factor not revealed by previously applied methods. Accordingly, linear methods as well as embedding algorithms based on linear optimization problems, such as PCA and Vivaldi, are insensitive to subtle structural features of the network, thereby explaining the disappointing degree of accuracy observed in previous studies. Furthermore, our discovery of a dimensionality-reducing decomposition of the delay space lends support to the theory that embedding techniques based on hierarchical decompositions [46] may outperform existing techniques that attempt to embed the entire distance matrix in one shot.

We leave for future work the applicability of nonlinear dimensionality reduction techniques originally developed in the machine learning and pattern recognition communities, e.g., [20, 41]. These methods, including the Isomap method briefly studied in Section 5.2 as well as more recent diffusion-based techniques [20], are based on metric embeddings which can reflect the intrinsic geometry of datasets and allow multiscale analysis for solving dimensionality reduction, clustering and parametrization.

Finally, we expect that the future availability of more comprehensive datasets, containing a larger number of representatives, including hosts in edge networks, and path information (i.e., traceroutes) will allow us to better quantify the effectiveness of the fractal measures and discover subtle properties of the delay space geometry that cannot be fully contemplated via analyzing the King datasets.

Acknowledgments

The authors would like to thank Emin Gün Sirer and Bernard Wong for making their raw measurements available, and Mark Crovella, Paul Francis, and Bruce Maggs for providing fruitful comments on this research. Finally, we are grateful to Niel Spring for his service as shepherd to this publication.

7. REFERENCES

[1] Hostip.info database. http://www.hostip.info.
[2] P2PSim. http://www.pdos.lcs.mit.edu/p2psim/.
[3] University of Oregon RouteViews Project. http://www.antc.uoregon.edu/route-views/.
[4] S. Abe and N. Suzuki. Omori's law in the Internet traffic. *Europhysics Letters*, 61(6), 2003.
[5] I. Abraham, Y. Bartal, H. T.-H. Chan, K. Dhamdere, A. Gupta, J. Kleinberg, O. Neiman, and A. Slivkins. Metric embeddings with relaxed guarantees. In *Proc. of IEEE FOCS*, 2005.
[6] R. E. Bellman. *Adaptive Control Processes*. Princeton University Press, 1961.
[7] A. Belussi and C. Faloutsos. Estimating the selectivity of spatial queries using the correlation fractal dimension. In *Proc. of VLDB*, 1995.
[8] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel J. Math.*, 52, 1985.
[9] M. Bădoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proc. of ACM STOC*, 2005.
[10] M. Costa, M. Castro, A. Rowstron, and P. Key. PIC: Practical Internet coordinates for distance estimation. In *Proc. of ICDCS*, 2004.
[11] M. E. Crovella and A. Bestavros. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6), 1997.

[12] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: a decentralized network coordinate system. In *Proc. of ACM SIGCOMM*, 2004.
[13] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. *SIGCOMM Comput. Commun. Rev.*, 29(4), 1999.
[14] K. P. Gummadi, S. Saroiu, and S. D. Gribble. King: estimating latency between arbitrary Internet end hosts. *SIGCOMM Comput. Commun. Rev.*, 32(3), 2002.
[15] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proc. of IEEE FOCS*, 2003.
[16] B. Huffaker, M. Fomenkov, D. J. Plummer, D. Moore, and K. Claffy. Distance metrics in the Internet. In *Proc. of IEEE Intl. Telecom. Symposium (ITS)*, 2002.
[17] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. In *Proc. of FOCS*, 2001.
[18] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Math.*, 26, 1984.
[19] J. Kleinberg, A. Slivkins, and T. Wexler. Triangulation and embedding using small sets of beacons. In *Proc. of IEEE FOCS*, 2004.
[20] S. Lafon and A. B. Lee. Diffusion maps and coarse-graining: A unified framework for dimensionality reduction, graph partitioning, and data set parameterization. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 28(9):1393–1403, 2006.
[21] E. Lebhar, P. Fraigniaud, and L. Viennot. The inframetric model for the Internet. In *Proc. of IEEE INFOCOM*, 2008.
[22] J. Ledlie, P. Gardner, and M. Seltzer. Network coordinates in the wild. In *Proc. of USENIX NSDI*, 2007.
[23] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15, 1990.
[24] E. K. Lua, T. Griffin, M. Pias, H. Zheng, and J. Crowcroft. On the accuracy of embeddings for Internet coordinate systems. In *Proc. of ACM/SIGCOMM Internet Measurement Conference*, 2005.
[25] H. V. Madhyastha, T. Anderson, A. Krishnamurthy, N. Spring, and A. Venkataramani. A structural approach to latency prediction. In *Proc. of ACM/SIGCOMM Internet Measurement Conference*, 2006.
[26] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani. iplane: An information plane for distributed services. In *Proc. of OSDI*, 2006.
[27] K. R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf. An introduction to kernel-based learning algorithms. *IEEE Transactions on Neural Networks*, 12(2), 2001.
[28] T. S. E. Ng and H. Zhang. Predicting Internet network distance with coordinates-based approaches. In *Proc. of IEEE INFOCOM*, 2002.
[29] R. V. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. In search of the elusive ground truth: the internet's as-level connectivity structure. In *SIGMETRICS*, 2008.
[30] V. Paxson and S. Floyd. Wide-area traffic: The failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3), 1995.
[31] V. Paxson and S. Floyd. Why we don't know how to simulate the Internet. In *Proc. of Winter Simulation Conference*, 1997.
[32] M. Pias, J. Crowcroft, S. Wilbur, S. Bhatti, and T. Harris. Lighthouses for scalable distributed location. In *Proc. of IPTPS*, 2003.
[33] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Topologically aware overlay construction and server selection. In *Proc. of IEEE INFOCOM*, 2002.
[34] M. Schroeder. *Fractal, Chaos and Power Laws: Minutes from an Infinite Paradise*. W. H. Freeman and Co., NY, 1990.
[35] Y. Shavitt and T. Tankel. Big-bang simulation for embedding network distances in Euclidean space. *IEEE/ACM Transactions on Networking*, 12(6), 2004.
[36] Y. Shavitt and T. Tankel. On the curvature of the Internet and its usage for overlay construction and distance estimation. In *Proc. of IEEE INFOCOM*, 2004.
[37] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *Proc. of ACM SIGCOMM*, 2003.
[38] G. Strang. *Linear Algebra and its Application, 2nd. Ed.* Academic Press, 1980.
[39] L. Tang and M. Crovella. Virtual landmarks for the Internet. In *Proc. of ACM/SIGCOMM Internet Measurement Conference 2003*, 2003.
[40] L. Tang and M. Crovella. Geometric exploration of the landmark selection problem. In *Proc. of Passive and Active Measurement Workshop*, 2004.
[41] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500), 2000.
[42] G. Wang, B. Zhang, and T. S. E. Ng. Towards network triangle inequality violation aware distributed systems. In *Proc. of ACM/SIGCOMM Internet Measurement Conference*, 2007.
[43] W. Willinger, W. E. Leland, M. S. Taqqu, and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1), 1994.
[44] B. Wong, A. Slivkins, and E. G. Sirer. Meridian: a lightweight network location service without virtual coordinates. *SIGCOMM Comput. Commun. Rev.*, 35(4), 2005.
[45] B. Zhang, T. S. E. Ng, A. Nandi, R. Riedi, P. Druschel, and G. Wang. Measurement based analysis, modeling, and synthesis of the Internet delay space. In *Proc. of ACM/SIGCOMM Internet Measurement Conference*, 2006.
[46] R. Zhang, C. Hu, X. Lin, and S. Fahmy. A hierarchical approach to Internet distance prediction. In *Proc. of IEEE International Conference on Distributed Computing Systems*, 2006.