

Compact Routing on Internet-like Graphs

Dmitri Krioukov *

Kevin Fall †

Xiaowei Yang ‡

February 2, 2008

Abstract

The Thorup-Zwick (TZ) routing scheme is the first generic stretch-3 routing scheme delivering a nearly optimal local memory upper bound. Using both direct analysis and simulation, we calculate the stretch distribution of this routing scheme on random graphs with power-law node degree distributions, $P_k \sim k^{-\gamma}$. We find that the average stretch is very low and virtually independent of γ . In particular, for the Internet interdomain graph, $\gamma \sim 2.1$, the average stretch is around 1.1, with up to 70% of paths being shortest. As the network grows, the average stretch slowly decreases. The routing table is very small, too. It is well below its upper bounds, and its size is around 50 records for 10^4 -node networks. Furthermore, we find that both the average shortest path length (i.e. distance) \bar{d} and width of the distance distribution σ observed in the real Internet inter-AS graph have values that are very close to the minimums of the average stretch in the \bar{d} - and σ -directions. This leads us to the discovery of a unique critical quasi-stationary point of the average TZ stretch as a function of \bar{d} and σ . The Internet distance distribution is located in a close neighborhood of this point. This observation suggests the analytical structure of the average stretch function may be an indirect indicator of some hidden optimization criteria influencing the Internet's interdomain topology evolution.

1 Introduction

The recent observations of and scalability concerns with the dynamics of the BGP routing table size growth, [9, 50, 52, 51, 25, 14], bring up the question of how small the routing table sizes for distributed routing on realistic massive graphs can be made *in principle*. In other words, what are the *fundamental* limits of compactness of graph-theoretic routing in such networks?

Answering this question involves two things. On one hand, it calls for assessment of results obtained in the area of distributed routing. Since our first interest is the lower limits that can be achieved *in principle*, we are more concerned with idealized *static* routing in this paper. Being the simplest and most fundamental routing model, static routing is where such limits can manifest themselves. In the more complicated dynamic case, these limits can only be higher.

On the other hand, answering the above question also requires understanding of the basic properties of *realistic* massive growing networks, the Internet being a good example of those. The structure and evolution of these networks are subjects of intensive studies these days.

1.1 Previous work

The above two research fields virtually do not intersect. We are not aware of any routing schemes designed specifically for scale-free graphs, and, vice versa, the literature concerned with the properties of scale-free nets has not addressed the routing problem yet. We tend to explain this lack of overlapping by the fact that understanding of the nature of large networks observed in reality started just recently. We review the results relevant to our work in the two separate subsections.

*dima@krioukov.net

†kfall@intel-research.net

‡yxw@mit.edu

1.1.1 Routing

Since the pioneering work by Kleinrock and Kamoun, [55], the trade-off between stretch¹ and the amount of routing information (i.e. memory space) required by a routing scheme has been understood, analyzed, and improved. Many relatively recent “Internet routing architecture” proposals, [16, 54], are based on the ideas of [55]. Since [55] was the first work of its type, it is not surprising that naïve routing based on it would generate stretch that is far from optimal. Indeed, the simple calculations presented in Appendix A show that the stretch produced by [55] on the present Internet interdomain graph would be of the order of 10.

However, the high value of stretch was not the central problem with [55]. This work along with the works that closely followed it, [56, 77], concentrated on optimal hierarchical network clustering (splitting into areas) satisfying a set of assumptions, but no proof of existence of such clustering for generic graphs and no algorithm to find it when it exists were obtained. While several subsequent works tried to overcome these problems, all known hierarchical routing schemes eventually turned out to be inferior with respect to more recent *direct* routing schemes.²

In the work by Peleg and Upfal, [76], the trade-off between memory space and stretch for *generic*³ networks was rigorously analyzed for the first time. The work contained several issues that were addressed in many publications that followed [76]. One of the issues was that only the total (per network) space was bounded, the local (per node) space was unbounded. Another issue was that the scheme required relabelling of nodes.

The fact that useful information about network topology can be embedded in node labels to reduce the space is most easily seen in the case of ring networks, where the local lower bound for shortest path routing is $\Omega(n)$ if relabelling is not allowed, and $\Omega(1)$ if nodes on the ring can be labelled sequentially.

The most efficient label set for an n -node network is obviously $[1, n]$.⁴ The fundamental lower bound is obtained in [33]. It is shown there that if the label space is $[1, n]$, then, for *any* stretch (including shortest path routing), there cannot exist a loop-free generic routing scheme that would guarantee the local space less than $\sim 3.7n^{1/2}$.

For *shortest path* routing, the lower bound turns out to be higher. The pessimistic but intuitively expected results obtained in [47] show that for *any* stretch-1 routing scheme, there exists a graph with maximum node degree d , $\forall d \in [3, n)$, such that $\Omega(n \log d)$ bits of memory are required at $\Theta(n)$ nodes. Since the trivial upper bound for shortest path routing⁵ is also $O(n \log d)$, [47] effectively demonstrates *incompressibility* of generic shortest path routing.

Fortunately, the majority of graphs are slightly better. Applying the Kolmogorov complexity theory to routing, the authors of [15] provide many upper and lower bounds for *almost all* graphs. In particular, for shortest path routing, not more than $3n$ (but not less than $n/2$) bits per node are shown to be enough for the $1 - 1/n^3$ portion of all $[1, n]$ -labelled graphs. If labelling is relaxed to allow for $O(\log^2 n)$ -sized labels, then the local memory space upper bound is reduced to $O(\log^2 n)$. The question about the space lower bound for stretch-1 routing on almost all graphs with free relabelling remains open.

While the results from the complexity theory or classical random graph theory ([39, 46]) for almost all graphs may induce some optimism, very little can be said, on the practical side, about if all the graphs from a given class of graphs are good or bad with respect to routing compactness. Exploring specifics of various graph families, a number of compact routing schemes for special types of graphs have been constructed. There are several results for rings, complete networks, trees, grids, decomposable, planar⁶, outplanar⁷, bounded genus⁸, chordal⁹, etc., graphs,—but none so far for scale-free graphs (even in the “almost all” context). It is easy to explain since the properties of scale-free networks from the graph-theoretic perspective are not fully understood yet.

Thus, the only currently existing tool to analyze the limits of compactness of routing on scale-free networks is generic routing schemes. Generic shortest path routing is incompressible, which means that if the memory space is to be reduced, then the stretch must be increased.

¹The stretch factor is a (usually worst-case) ratio of the path length produced by a routing scheme to the shortest path length. Stretch-1 routing and shortest path routing are synonymous.

²A routing scheme is *direct* if the output port calculated at every node depends on the destination label and nothing else. This implies that message headers cannot be altered by intermediate nodes, and, hence, paths produced by a direct routing scheme cannot have loops, which justifies word “direct.” Many hierarchical routing schemes are not direct.

³That is, *all*. A generic routing scheme is applicable to all graphs.

⁴The $[1, n]$ label set is relevant to a specific terminology used in the routing literature. The common term “routing table” denotes a direct routing scheme with labels from this set.

⁵The outgoing port is listed for every destination.

⁶A planar graph can be drawn on a plane without crossing edges.

⁷An outplanar graph can be embedded in a plane with all its vertices lying on a convex polygon.

⁸The genus of a graph is the minimum number of edge crossings with which a graph can be drawn on a plane. Planar graphs have genus 0.

⁹The graph is chordal if all its cycles longer than 4 have chords, or, equivalently, if it does not have induced cycles longer than 3.

Table 1: Total space lower bounds for low stretch values.

Stretch s	Lower bound	Source
$1 \leq s < 1.4$	$\Omega(n^2 \log n)$	[47]
$1.4 \leq s < 3$	$\Omega(n^2)$	[44]
$3 \leq s < 5$	$\Omega(n)$	[87]

The memory space lower bound dependence on stretch is not “continuous.” As shown in [47], any generic routing scheme with the maximum stretch strictly less than 1.4 must use at least $\Omega(n \log n)$ bits of memory on some nodes of some graphs. In other words, the lower bound for generic schemes with stretch $s < 1.4$ is the same as in the incompressible case of shortest path routing (consider the bound of $\Omega(n \log d)$ discussed above and take $d = \Theta(n)$). Furthermore, as shown in [44], the lower bound for schemes with stretch strictly less than 3 is nearly the same as for shortest path routing— $\Omega(n)$ bits of memory on some nodes of some graphs.

The minimum stretch factor that allows for significant memory space lower bound decrease is 3. Cowen introduces a very simple direct stretch-3 routing scheme with the local memory space upper bound of $O(n^{2/3} \log^{4/3} n)$ in [22]. The scheme uses relabelling, and labels are of size $3 \log n$. In [88], Thorup and Zwick improve Cowen’s result and deliver a local space upper bound of $O(n^{1/2} \log^{1/2} n)$. They also show how to implement routing decisions at *constant* time per node and to reduce the label size to $(1 + o(1)) \log n$. We call the above two stretch-3 schemes the *Cowen* and *Thorup-Zwick* (TZ) schemes respectively.

The local memory space upper bound provided by the TZ scheme is nearly (up to a logarithmic factor) optimal (best possible) since, as demonstrated in [87], any generic routing scheme with stretch strictly less than 5 must use at least $\Omega(n^{1/2})$ bits of memory on some nodes of some graphs (see Table 1). To the best of our knowledge, the TZ scheme is the single generic stretch-3 routing scheme delivering a nearly optimal local memory upper bound today. This makes it “exceptional” in a sense that it delivers a nearly optimal first possibility to decrease the local space down from the shortest path routing incompressible limits. This also explains why it is a primary subject of our present work.

To finish the introduction to relevant results in routing, we briefly touch on two more areas.

First, nothing confines one to considering only a multiplicative stretch factor. The concepts of *additive* stretch—that is, the additive error factor in distance approximation—and even of mixed multiplicative-additive stretch have been introduced and studied to some degree ([3, 24, 34]). Additive stretch models are potentially better suited for graphs with low average distances (the Internet interdomain graph, for example) since, as can be easily seen, the *short* distances are harder to approximate than the long ones ([3, 22]¹⁰). However, there are very few *routing schemes* based on additive stretch. The latest one is probably a very efficient additive stretch-2 routing scheme for chordal graphs by Dourisboure and Gavaille, [32].

Second, the above discussion concerns the static case only. In case of dynamic networks, the other components that add complexity to the picture are stability issues as well as adaptation or communication costs¹¹.

One of the first works that rigorously addresses the problem of routing in dynamic networks is [1]. The “positive” result in the paper is a dynamic routing scheme for growing trees. One of the interesting “negative” results is the analysis of the trade-off between the stretch and adaptation cost. Assuming that the space and message sizes are unbounded, [1] shows that any generic dynamic routing scheme with stretch $s < k$ must send at least $\Omega(n/k)$ messages per topology change on some networks.

In a recent work by Krizanc, Luccio, and Raman, [58], three schemes for dynamic routing on rings with different stretch-space-adaptation trade-offs are constructed. All of the three schemes are dynamic versions of interval routing schemes.¹²

The *Bubbles* model, [23], is a *generic* dynamic routing scheme that uses hierarchical partitioning of the spanning tree of a graph. Because of the specifics of the network model considered in [23], the stretch factor is replaced there by the “super-hop count,” the maximum number of hops produced by the scheme, and the adaptation cost is measured by “adaptability,” the maximum number of nodes affected by topology change updates. The most efficient variant of the scheme, designed for high-degree networks, provides the local memory space upper bound of $O(kn^{1+1/k} \log d)$, where k is the maximum number of hops and d is the maximum node degree. The adaptation cost upper bound is

¹⁰In fact, one of the very important component of many routing schemes including the Cowen scheme is a careful balance between short and long paths.

¹¹That is, the number of messages generated per topology change, their sizes, or the total amount of data sent to guarantee communication, etc.

¹²A review of interval routing is given in [42].

$O(3^k n^{1/k} d)$ with node failures and $O(3^k n^{1/k})$ with link-only failures. The adaptation cost lower bound for low-degree graphs ($d = O(1)$) is also obtained. It is $\Omega(n^{1/k})$.

Finally, in the context of the *end-to-end communication* problem (see review by Fich [38]), the stretch factor is not explicitly considered—the problem is just to guarantee communication between two fixed nodes in presence of frequent network failures and to optimize the trade-off between the total number of messages generated per data item (communication cost) and required local memory space per incident link. For the first solution that is polynomial in communication cost and logarithmic in memory space, see the work by Kushilevitz, Ostrovsky, and Rosén [60]. For the latest memoryless network result, see [40].

For more details on progress in routing, see the excellent review by Gavaille [43] and the monograph by Peleg [75].

1.1.2 Scale-free networks

Since the discovery of *power-law* distributions in the Internet in [37], the Internet *scale-free* nature has been a subject of very intense studies and generated an enormous number of publications. A particularly interesting fact is that the Internet appears to be just one example of scale-free networks that have been found extremely ubiquitous. The list of networks, within which power-law or, more generally, *fat-tailed* distributions have been observed, include—beyond the Internet, both at the interdomain and router levels ([89, 82])—the WWW ([5]), power grids ([6]), airport networks ([6]), biological ([53]), ecological ([69]), language ([28]), and social ([73]) networks, the latter including scientific collaboration ([70]), movie actor collaboration ([6]), human acquaintance ([6]), and sexual contact ([63]) networks.¹³

All the networks listed above involve an element of randomness. Classical Erdős-Rényi random n -node graphs, [35], have links between every pair of vertices with the uniform probability p . The ensemble of such graphs is called $\mathcal{G}_{n,p}$. Their average node degree is $\bar{k} \sim np$, the node degree distribution is the Poisson distribution with exponentially small number of high-degree nodes, $P_k \sim \bar{k}^k e^{-\bar{k}}/k!$, and average distance is $\bar{d} \sim \log n / \log \bar{k}$, [10].

All the networks observed in reality defer drastically from the $\mathcal{G}_{n,p}$ graphs. One of the differences of particular importance for this paper can be seen as certain inconsistency between the average distance and average node degree predicted by the Erdős-Rényi model. In the real Internet interdomain 1.1×10^4 -node graph, for example, $\bar{k} \sim 5.7$ and $\bar{d} \sim 3.6$, [89]¹⁴, while the $\mathcal{G}_{1.1 \times 10^4, 5.2 \times 10^{-4}}$ graphs have $\bar{d} \sim 5.3$. The $\mathcal{G}_{n,p}$ graphs of the same size with the right average distance $\bar{d} \sim 3.6$ would have to have the average degree $\bar{k} \sim 14$. In Sections 2.2.3 and 3, we see how strongly these slight differences affect the average stretch.

The simultaneously small values of the average distance and average node degree necessarily imply a larger portion of high-degree nodes than in the classical random graphs. In other words, the node degree distribution must be fat-tailed. The power-law, $P_k \sim k^{-\gamma}$, one of such fat-tailed distributions, is what has been observed in many networks listed above, exponent γ ranging between 2 and 3. For the Internet interdomain graph, $\gamma \sim 2.1$, [37, 89].

Both the $\mathcal{G}_{n,p}$ graphs and graphs with fat-tailed degree distributions are often said to possess the *small-world* property, [66], to emphasize that they have extremely low average distances (for networks of such size), even though average distances in $\mathcal{G}_{n,p}$ graphs are slightly higher. The famous play *Six Degrees of Separation*, [49], is based on the observation made in [66] that the average distance in human acquaintance networks is around 6. As far as routing in the Internet is concerned, the simple but critically important fact that the Internet distance distribution has very low values of mean and dispersion (that is, that there are virtually no remote points) gets fairly often either overlooked or neglected.

Networks with fat-tailed degree distributions are also called *scale-free* since their node degree distribution lacks any characteristic scale, [8], in contrast to the $\mathcal{G}_{n,p}$ graphs with the narrow Poisson degree distribution centered around the characteristic average value $\bar{k} \sim np$.

The explanation of appearance of fat-tailed distributions in realistic networks is obviously a very important problem. While a large number of models generating power-laws have been suggested, arguably none of them so far captures the underlying principles of the Internet evolution and predicts the observed Internet topology well enough.

The most popular model is the model for growing networks with preferential attachment¹⁵ by Barabási and Albert (BA), [8]. The BA model is very simple, it does not have external parameters, which makes it attractive for a physicist, but, in its “pure” form, it predicts $\gamma = 3$. The model can be freely modified to produce other values of γ and even scaling behaviors deviating from power-laws, [4], but its applicability to the Internet evolution has been extensively criticized in [94, 18]. In particular, in [94], it is noted that the BA model and its derivatives are capable of reproducing what has been already measured but they fail to predict correctly anything new about the Internet

¹³See also [93, 92, 71] and [81].

¹⁴For more detailed measurements of the Internet topology, see [13].

¹⁵The probability for a new-coming node to attach to a target node already in the network is proportional to the target node degree.

topology, anything that has not been measured yet. As such, the BA model is not *explanatory* but *descriptive*, in the terminology of [94].

One of the most interesting models for the Internet evolution is analyzed in [36]. It is shown there that a simultaneous (trade-off) optimization of last-mile costs (geometrical distance) and average hop distance from the network “center” (the first node arrived) can lead to power-laws. In other words, power-laws can be a result of optimization of the trade-off between the physical link cost and average delay associated with the average path length in hops—in data networks, every hop is a source of queuing delay and packet loss. Although there is arguably no such optimization intentionally (in a controlled manner) happening in the real Internet, which is driven primarily by its economy outlined, for example, in [74] and modelled in [57, 17], the link, noted in [36], to the Mandelbrot language model, [65], maximizing language efficiency and resulting in power-laws, deserves some serious attention.

It is worth mentioning that the whole subject of the scale-free nature of the Internet has been doubted in [61], where it is shown that just traceroute-based measurement techniques¹⁶ may be solely responsible for the Internet *appearing* scale-free, and it may belong to the $\mathcal{G}_{n,p}$ class in reality. Indeed, it is easy to see that the farther a measured node is from *all* the measuring nodes (sources of traceroutes) on a graph, the smaller portion of the total number of links incident to the measured node can be detected by traceroutes. Although some possibility for the Internet *router level* topology to be less fat-tailed does exist, the results of analysis in [61] itself essentially rule out any possibility for the Internet *interdomain* topology to deviate strongly from the power-law if one considers the very low value of the almost undoubtedly measured average distance and the high enough numbers of measuring points used in various Internet interdomain topology studies (ten vantage points are used, for example, in [83]).

On the practical side, a very useful work is presented in [85]. The authors distill a set of criteria assigning a characteristic “signature” to any type of graphs. These signatures are used then to compare the real Internet topology with topologies produced by various Internet topology generators and with topologies of several standard types of graphs (trees, grids, complete graphs, classical random graphs, etc.). Surprisingly enough, the structural Internet topology generators trying to incorporate the perceived hierarchical structure of the Internet in their algorithms are found inferior to the degree-based generators “blindly” reproducing the observed degree distribution. The simplest generator of this type is the PLRG generator suggested in [2]¹⁷ and analyzed (among other generators) in [85]. It is also found in [85] that the only type of standard graphs having the same signature as the Internet is complete networks.

Another important observation made in [85] is about the presence of correlation between the link value, defined as a weighted number of shortest paths passing via the link, and the lower degree of the nodes attached to the link. This observation is consistent with the earlier measurements of the Internet interdomain graph in [90] showing that the node betweenness, defined as the total number of shortest paths passing via a node, is linearly correlated with the node degree.

These measurements point to the “self-establishing” nature of the Internet hierarchy, which is further revealed by the observations of the power-law decay of the clustering coefficient in [90, 89]. The clustering coefficient c_k , defined as the average ratio of the number of 3-cycles involving k -degree nodes to its maximum value $k(k-1)/2$, measures how close an average k -degree node neighborhood is to a clique. Its power-law decay for the Internet interdomain graph indicates that small, low-degree ASs¹⁸ tend to create numerous, highly clustered structures that are connected with each other via a sparse formation of “hubs”—large, high-degree ASs.

It is interesting to note that the power-law decay of the clustering coefficient has been observed only for “uncontrolled,” “self-evolving” networks—the Internet interdomain graph, the WWW, biological, language, and social networks. Networks with a stronger element of design and external control—the Internet router level graph and power grids, for example—do not exhibit such behavior, [89, 78, 79]. The clustering coefficient as a function of node degree seems to be relatively constant in the latter cases, while its average values are still much higher than in the $\mathcal{G}_{n,p}$ graphs, which is another drastic difference between the $\mathcal{G}_{n,p}$ model and real-world networks, [93, 71].

For the analytical part of our present work, we need to know the distance distribution in scale-free graphs. The problem is very hard and it has not been solved analytically yet. There are some recent results on the *average* distance in scale-free networks, [19, 21]. *Implicit* expressions for the distance *distribution* are constructed in [20]. More *explicit* analysis of the distance distribution is performed in [31, 41].¹⁹ Unfortunately, all these results are valid only for static, equilibrium networks without vertex-vertex degree correlations. All realistic scale-free networks are growing, non-equilibrium. They necessarily have node degree correlations resulting in much wider distance distributions, [26]. Surprisingly, the model constructed by Dorogovtsev, Goltsev, and Mendes for deterministic scale-free graphs in [27]

¹⁶These obviously include both standard traceroutes and BGP table dumps.

¹⁷The construction procedure is due to Molloy and Reed, [67, 68].

¹⁸For the measurements of correlation between AS size and degree, see [84].

¹⁹Why the frequently referred expressions derived in [72] are imprecise is shown, for example, in [59].

(the DGM model) turns out to be capable of analytically producing a Gaussian distance distribution similar to the distance distribution observed in the real Internet, [89, 90, 14, 9]. The width of the Gaussian for a 10^4 -node network, 1.1, is very close to the width of the Internet interdomain distance distribution, 0.9, but the average distance is slightly higher—4.8 instead of 3.6. As noted in [27], simulation-based measurements of the distance distribution in the BA model also produce similar Gaussians, [59].

For further details on scale-free networks, see the excellent review [29] and book [30] by Dorogovtsev and Mendes.

1.2 Our contribution

One might expect that for scale-free graphs, the majority of known generic routing schemes would be very inefficient. Indeed, many routing schemes (including the Cowen and TZ schemes) incorporate *locality* by carefully differentiating between close and remote nodes. This approach makes routing more efficient (in the stretch-versus-space trade-off sense) by keeping only approximate (non-shortest path) routing information for remote nodes, while full (shortest path) routing information is kept for local nodes. In scale-free graphs characterized by low average distances and distance distribution widths, local nodes comprise huge portions of all the nodes in a network, so that one might suspect that locality-sensitive approaches might break for such networks. For a good example demonstrating that this might be quite plausible, see the Appendix A, where the stretch factor is found to be very high for the Kleinrock-Kamoun (KK) routing scheme [55] applied to the scale-free networks.

Furthermore, one can take the situation to its extreme and consider a “smallest-world” graph, that is, a complete graph. The idea is suggested in part by [85], where the Internet graph “signature” is found to be similar to the complete network “signature” (cf. Section 1.1.2). One would find then that both the Cowen and TZ average stretch factors in this extreme case of complete graphs are high—as can be easily checked, the average TZ stretch for a complete graph of size n is $2 - n^{-1/2} \log^{-1/2} n - o(n^{-1/2})$,²⁰

We find that the case of realistic scale-free networks with Internet-like characteristics is significantly better.

We consider the TZ scheme, which is an “exceptional” routing scheme in the sense explained in Section 1.1.1. Being generic, the TZ scheme provides only general maximum stretch and space bounds. It says nothing about the average stretch or stretch distribution on a particular class of graphs.

We calculate, both analytically and via simulations, the TZ stretch distribution on Internet-like topologies. The analytical part of the problem is hard. It assumes knowledge of the distance distribution in correlated scale-free networks. The exact form of this distribution has not been obtained analytically yet (see Section 1.1.2). Given the observation that the DGM model [27] analytically produces the Gaussian distance distribution that is close to the real Internet distance distribution, we choose to parameterize distance distributions in small-world graphs we consider in this paper by Gaussian distributions. To obtain our results we still have to make a series of simplifying assumptions that are fully discussed in Section 2.1.

For the simulation part, we develop our own TZ scheme simulator and use it on graphs produced by our implementation of the PLRG generator [2], the initial justifications for using it being discussed in Section 1.1.2. Since the PLRG generator outputs uncorrelated networks, there are some concerns regarding its capability of reproducing *all* the features of strongly correlated nets, such as the Internet. However, since, as we see in Section 2.1, the stretch distribution turns out to be a function of the distance distribution and the graph size only, all we need from a graph generator for our purposes is that distance distributions in graphs produced by it be close to distance distributions observed in real-world graphs. We find that PLRG-generated graphs with the node degree distribution exponent $\gamma = 2.1$ have the distance distribution that is very close to the distance distribution observed in the Internet.

We obtain a close match between the analysis and simulation data for the average TZ stretch and stretch distribution in Section 2.2. We find that the average stretch is *very low* and virtually *independent* of exponent γ . In particular, in the case of the Internet interdomain graph, $\gamma \sim 2.1$ and size $n \sim 10^4$, the average stretch is 1.14 and 1.09 according to the analysis and simulations respectively. The stretch distribution has a peculiar form. The majority of paths produced by the TZ scheme are shortest—up to 71% according to the simulations. The majority of non-shortest paths have stretch values of $4/3$ and $5/4$. The portion of paths with other stretch values is very small.

The average number of entries in the routing table²¹ is also extremely low—well below its upper bounds. For graphs with the Internet-like parameters, $n \sim 10^4$, $\gamma \sim 2.1$, it is approximately 52.

We also show that the average stretch slowly *decreases* with the network growth even if the average distance scales as $\log n$. However, the average stretch does not approach 1 even for sufficiently large n . Therefore, the amount of non-shortest paths seems to be unbounded.

²⁰The Kleinrock-Kamoun average stretch is much worse, of course. It is trivial to see that it grows as $\Theta(\log n)$.

²¹We are still using term “routing table” here even though it is not completely correct from the graph-theoretical perspective since the TZ scheme does not use labels from the $[1, n]$ set.

The fact that the average stretch on scale-free networks turns out to be low is not by itself surprising. Indeed, a scale-free graph can be described as a “collection of interconnected stars,” and it is easy to see that both the Cowen and TZ average stretch factors on stars are equal to 1. The average stretch is expected to be substantially higher for other types of random networks. We confirm this expectation in Section 2.2.3 where we calculate the average stretch for certain $\mathcal{G}_{n,p}$ graphs.

We also obtain a row of really surprising results presented in Section 3. The analytical expressions we provide for the average stretch \bar{s} allow us to consider it as a function of the parameters of the distance distribution in a graph, the parameters being the average distance \bar{d} (the first moment) and distance distribution width σ (the square root of the second moment or the standard deviation). First, we find that both \bar{d} and σ of the distance distribution in the Internet are very close to the local *minimums* of $\bar{s}(\bar{d}, \sigma)$ in the \bar{d} - and σ -directions respectively.

Next, simultaneous proximity of the Internet distance distribution to the minimums of $\bar{s}(\bar{d}, \sigma)$ in the both directions makes us search for a stationary point²² and potential extremum of \bar{s} . Our analytical results allow us to collect enough data to discover: 1) a region of \bar{d} and σ , where function $\bar{s}(\bar{d}, \sigma)$ is concave and stretch is particularly low, which we call the *minimal stretch region* or the MSR; and 2) a unique critical *quasi-stationary* point of \bar{s} at the edge of the MSR, which we call the MSR *apex*. The apex is characterized by the shortest distance between the sets of minimums of \bar{s} in the \bar{d} - and σ -directions. The two sets do not intersect but are extremely close to each other at the apex. The surface of the average stretch function values in the apex neighborhood consists solely of elliptic points,²³ but the minimal deformation of the surface towards the potential intersection appears to result in a unique parabolic point.

The points corresponding to distance distributions of all random graphs with power-law node degree distributions lie in the MSR. In addition, the Internet distance distribution is located in a very *close neighborhood* to the MSR apex. Even a stronger statement is valid: $\gamma = 2.1$ is the value of γ corresponding to the distance distribution that is *closest* to the apex, compared to all other values of γ . The $\mathcal{G}_{n,p}$ graphs are far away both from the MSR and from its apex.

The phenomena outlined above appear to be a reflection of existence of a certain link between the Internet topology and the analytical structure of the average TZ stretch function. This is quite unexpected since the Internet, as we know it today, seems to have nothing to do with stretch, in general, and with the TZ stretch, in particular. That is why these effects cannot be fully interpreted within the set of ideas we operate with in this paper. Although, see Section 4 for some hints towards possible explanations.

In a recent work dedicated to a specific subject, [45], Gavaille and Nehéz raise a very important general issue of application of results in “theoretical” routing to *realistic* networks.²⁴ To the best of our knowledge, our work is among the first ones trying to create a link between routing and realistic scale-free networks. The principal result of this paper showing that the TZ stretch on Internet-like graphs is low, opens a well-defined path for the future work in this area, as further discussed in Sections 4 and 5.

2 Stretch

Both the Cowen and TZ schemes are very simple. They involve four separate components: the landmark set (LS) construction procedure, routing table construction, labelling, and routing itself. The TZ scheme differs from the Cowen scheme by improving just the first part; the other three are the same. We remind the outline of the TZ scheme below.

The scheme operates on any undirected graph $G = (V, E)$ with positive edge weights. Let $n = |V|$ be the graph size, $\delta(u, v)$ be the distance between a pair of nodes $u, v \in V$, L be the LS, $L(v)$ be a landmark node closest to node $v \in V$, and $C(v)$ be v 's cluster defined for $\forall v \in V$ as a set of all nodes c that are closer to v than to their closest landmarks,

$$C(v) = \{ c \in V \mid \delta(c, v) < \delta(c, L(c)) \}. \quad (1)$$

²²The point is stationary if all first-order partial derivatives of a function at this point are zero. This is a necessary condition for an extremum. The function has a minimum (maximum) at a given stationary point if all eigenvalues of the matrix of all second-order partial derivatives of the function at this point are positive (negative).

²³Any point on a regular surface is always of one of the following four types: planar (example: any point on a plane), elliptic (examples: any point on an ellipsoid, peaks and pits), parabolic (examples: any point on a cylinder, ridges and channels), and hyperbolic (examples: any point on a hyperboloid, passes). A function of two arguments has an extremum at its stationary point if the corresponding point on a surface of its values is elliptic.

²⁴They also question what realistic networks are. We believe that this question is being actively answered in the work discussed in Section 1.1.2.

Clusters are similar to the Voronoi diagrams but they can intersect. If $l \in L$, then $L(l) = l$ and $C(l) = \emptyset$ by definition. If L is empty, then for $\forall v \in V$, $L(v) = \emptyset$ and $C(v) = V$.

The TZ LS construction algorithm interactively selects landmarks from the set of large-cluster nodes W . At the first iteration, $W = V$ and every node $w \in W$ is selected to be a landmark with a specific uniform probability q/n with $q = (n/\log n)^{1/2}$. The expected LS size after the first iteration is q . At the subsequent iterations, W is redefined to be a set of nodes that have clusters of size greater than a specific threshold $\tilde{q} = 4n/q$,

$$W = \{ w \in V \mid |C(w)| > \tilde{q} \}, \quad (2)$$

and additional portions of landmarks are selected from W with a uniform probability $q/|W|$. The iterations proceed until W is empty.

Every node $v \in V$ calculates then its outgoing port for the shortest path to every $l \in L$ and every $c \in C(v)$. This is the routing information that is stored locally at v . As one can see, the essence of the LS construction procedure is the right balance between the LS and cluster sizes (or, effectively, between q and \tilde{q}). The cluster sizes are upper-bounded by definition (2), and the involved part of the proof is to demonstrate that the algorithm terminates with a proper limit for the expected LS size, which turns out to be $2q \log n$. This guarantees the overall local memory upper bound of $O(n^{1/2} \log^{1/2} n)$.

The label of node v (used as its destination address in packet headers) is then a triple of its ID, the ID of its closest landmark $L(v)$, and the local ID of the port at $L(v)$ on the shortest path from $L(v)$ to v . With these labels, routing of a packet destined to v at some (intermediate) node u occurs as follows: if $v = u$, done; if $v \in L \cup C(u)$, the outgoing port can be found in the local routing table at u ; if $u = L(v)$, the outgoing port is in the destination label in the packet; otherwise, the outgoing port for the packet is the outgoing port to $L(v)$ —the $L(v)$ ID is in the label and the outgoing port for it can be found in the local routing table. The demonstrations of correctness of the algorithm and that the maximum stretch is 3 are straightforward ([22, 88]).

2.1 Analytical results

In this section, we provide analytical expressions for the TZ stretch distribution on a small-world graph with a given distance distribution, in general, and with the Gaussian distance distribution, in particular.

We start with the following assumption drastically simplifying the analysis:

Assumption 1 *Only the first iteration of the LS construction algorithm is considered.*

There are two justifications making this assumption reasonable. First, as shown both in [87] and below in Claim 3, the first iteration guarantees that the *average* cluster size is below n/q ; the subsequent iterations guarantee that *all* cluster sizes are upper-bounded by $4n/q$. Therefore, the error introduced by this assumption for the *average* stretch is small as we see in the next section. The second justification making the error particularly small is that we consider small-world graphs which have very short average distances and narrow distance distributions. Indeed, if there are no long distances in a graph, then even after just the first iteration, the majority of clusters are small. The error introduced by the assumption is related to the difference between the expected LS size q after the first iteration and the average LS size observed in simulations, which is reported in the next section.

The fact that ratio between the expected LS size and the graph size is infinitesimally small for large graphs, $q/n \xrightarrow{n \rightarrow \infty} 0$, makes the following claim true:

Claim 1 *The difference between the distance distribution in G and distance distribution in its subgraphs \tilde{G} induced by $V \setminus \tilde{L}$, $\forall \tilde{L} \subset L$, can be neglected for large graphs.*

For the rest of this section, we let q denote the actual size of the LS, $q = |L|$. We also denote the distance p.d.f. and c.d.f. by $f(d)$ and $F(d)$ respectively. With D being the graph diameter, we allow $d = 0 \dots D$, where $f(0) = 1/n$ is the probability of zero-distance from a random node to itself. In some places below, we also refer to the continuous limit approximation (that is, to the assumption that $f(d)$ is continuous), but we explicitly avoid using it in the evaluations of the next section.²⁵ With the above notations and Assumption 1, we are ready to formulate the following claim:

²⁵Note, however, that the discrete case can often be closely approximated by the continuous case. Indeed, recall that as soon as $f(d)$ is sufficiently smooth and a sufficient number of its first derivatives are small enough at the interval boundaries $[0, D]$, which is the case for the Gaussian form of $f(d)$ we eventually select, then, according to the Euler-Maclaurin sum formula, the sum becomes indistinguishable from the integral over the same interval, $\sum_{d=0}^D f(d) \rightarrow \int_0^D f(d) dd$.

Claim 2 *The p.d.f. for the distance between a random node and its i 'th closest landmark is given by*

$$g_i(d) = c_i F(d)^{i-1} f(d) (1 - F(d))^{q-i} \quad (3)$$

with normalization coefficients $c_i = i \binom{q}{i}$ in the continuous limit.

Expression (3) is intuitively expected since $F(d)^{i-1}$ approximates the probability that $i-1$ landmark nodes are closer than d , and $(1 - F(d))^{q-i}$ approximates the probability that the rest of landmark nodes are farther than d (cf. the *order statistics*, [80]). For a more rigorous proof of Claims 1 and 2 above, see Appendix B.1.

Given the p.d.f. for the distance to the closest landmark $g_1(d)$ in (3), one can easily prove (see Appendix B.2) the following claim:

Claim 3 *The average cluster size $|\overline{C}| \leq n/(q+1)$.²⁶*

We next denote by $g(d)$ the p.d.f. for the average distance to *all* landmarks,

$$g(d) = \frac{1}{q} \sum_{i=1}^q g_i(d). \quad (4)$$

Since landmarks are just some q random nodes, $g(d)$ is equivalent to $f(d)$. In the continuous limit,²⁷ we have

$$g(d) = f(d) \frac{1}{q} \sum_{i=1}^q i \binom{q}{i} F(d)^{i-1} (1 - F(d))^{q-i} = f(d). \quad (5)$$

Letting w be the source node and v be the destination, we fix the notation for the following three random variables:

$$x = \delta(w, L(v)), \quad \text{p.d.f.} = g(x), \quad (6)$$

$$y = \delta(v, L(v)), \quad \text{p.d.f.} = g_1(y), \quad (7)$$

$$z = \delta(w, v), \quad \text{p.d.f.} = f(z). \quad (8)$$

With these notations, the random variable for the approximate stretch is

$$s^*(x, y, z) = \frac{x + y}{z}. \quad (9)$$

This expression for stretch is approximate for two reasons. First, it does not account for stretch-1 paths to destinations in the local cluster. Second, it does not incorporate the *shortcut effect*. Recall that the Cowen routing algorithm is such that if destination $v \notin L$ and if a message on its way to $L(v)$ passes some node $u \mid v \in C(u)$, then the message never reaches $L(v)$ but goes along the shortest path from u to v . In Appendix B.3, we justify the following claim:

Claim 4 *The stretch-1 and shortcut paths can be approximated by the following correction to s^* in (9):*

$$s(x, y, z) = \begin{cases} 1 & \text{if } z < y, \\ 1 & \text{if } z < x, \\ \frac{x+y}{z} & \text{otherwise.} \end{cases} \quad (10)$$

Our problem now is to find the joint p.d.f. $t(x, y, z)$ for $s(x, y, z)$. If x , y , and z were independent random variables, then $t(x, y, z)$ would be given by $g(x)g_1(y)f(z)$. They are not independent by definitions (6)-(8), which result in the triangle inequality,

$$|x - y| \leq z \leq x + y. \quad (11)$$

Furthermore, there can be some other correlations in the distance matrix. To proceed, we make the following assumption:

Assumption 2 *There are no correlations in the distance matrix, other than those associated with the triangle inequality.*

²⁶Note that $q+1$ instead of q in the denominator guarantees the right bound for the case of the empty LS, $q=0$.

²⁷In the discrete case and with $f(d)$ being Gaussian, $g(d)$ is still virtually identical to $f(d)$ because of the Euler-Maclaurin sum formula.

With this assumption, we can prove (see Appendix B.4) the following claim:

Claim 5 *The stretch p.d.f. is given by*

$$t(x, y, z) = \frac{g(x)g_1(y)f_t(x, y, z)}{F(x+y) - F(|x-y|)}, \quad f_t(x, y, z) = \begin{cases} f(z) & \text{if } |x-y| \leq z \leq x+y, \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

This claim is intuitively expected—the triangle inequality (11) just cuts a corresponding portion out from $f(z)$ with the proper normalization coefficient.

The average stretch and the stretch distribution are now

$$\bar{s} = \sum_{x, y, z=0}^D s(x, y, z)t(x, y, z), \quad (13)$$

$$\rho(\varsigma) = \sum_{\substack{x, y, z=0 \\ s(x, y, z)=\varsigma}}^D t(x, y, z). \quad (14)$$

In the above expression for the stretch distribution $\rho(\varsigma)$, the summation is over such values of x , y , and z that their transformation according to (10) yields ς .

Equations (13) and (14) are our final analytical results that we require for the numerical evaluations of the next section. Of particular note is that the stretch distribution and average depend only on $f(d)$ and q .

At this point, however, we may try to substitute any specific form of the distance distribution into (13) and (14). As discussed in the introduction, we are interested in the Gaussian distance distribution with the average distance \bar{d} and standard deviation (width) σ ,

$$f(d) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{d-\bar{d}}{\sigma}\right)^2}. \quad (15)$$

Assuming that the distribution is continuous,²⁸ we can express the distance c.d.f. via the error function,

$$F(d) = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{d-\bar{d}}{\sigma\sqrt{2}} \right) \right]. \quad (16)$$

The average stretch becomes the following integral:

$$\bar{s} = \iiint s(x, y, z)t(x, y, z) dx dy dz, \quad (17)$$

which, after a series of substitutions, transforms to

$$\bar{s}(\bar{d}, \sigma) = \frac{2^{2-q}q}{\sigma^3(2\pi)^{3/2}} \iint \frac{e^{-\frac{1}{2\sigma^2}[(x-\bar{d})^2+(y-\bar{d})^2]} \left[1 - \operatorname{erf} \left(\frac{y-\bar{d}}{\sigma\sqrt{2}} \right) \right]^{q-1}}{\operatorname{erf} \left(\frac{x+y-\bar{d}}{\sigma\sqrt{2}} \right) - \operatorname{erf} \left(\frac{|x-y|-\bar{d}}{\sigma\sqrt{2}} \right)} dx dy \int_{|x-y|}^{x+y} e^{-\frac{1}{2}\left(\frac{z-\bar{d}}{\sigma}\right)^2} s(x, y, z) dz. \quad (18)$$

Unfortunately, we cannot evaluate even the inner-most integral in any special functions known either to us or to Gradshteyn-Ryzhik [48]. Therefore, we retreat to numerical evaluations of (13) and (14) in the explicitly discrete case.

2.2 Numerical results and simulations

In our numerical evaluations of (13) and (14), x , y , and z (defined in (6)-(8)) are integer variables with the following ranges:

$$x, y = 1 \dots D, \quad (19)$$

$$z = \max(1, |x-y|) \dots \min(D, x+y), \quad (20)$$

$$D = \lceil \bar{d} \rceil + \lceil 10\sigma\sqrt{2} \rceil, \quad (21)$$

²⁸In fact, the Gaussian distribution is continuous by definition. According to the de Moivre-Laplace theorem, it is an asymptotic form of the binomial distribution, $\binom{D}{d} \vartheta^d (1-\vartheta)^{D-d} \xrightarrow{D \rightarrow \infty} (\sigma\sqrt{2\pi})^{-1} e^{-(d-\bar{d})^2/(2\sigma^2)}$ with $\vartheta = \bar{d}/D$ and $\sigma^2 = D\vartheta(1-\vartheta)$. Although the values of $D \sim 13$, $\bar{d} \sim 3.6$, and $\sigma \sim 0.9$ observed in the Internet make this approximation essentially invalid for analytical purposes, we can still use (15) with discrete d for *numerical* computations.

Table 2: The top ten stretch values and percentage of paths associated with them.

Stretch	Analysis (%)	Simulations (%)
1	58.7	70.8
4/3	16.0	13.1
5/4	14.8	9.71
3/2	4.95	2.33
5/3	2.88	0.731
6/5	2.10	2.54
2	0.434	0.210
7/5	0.173	6.77×10^{-2}
7/6	5.20×10^{-2}	0.460
8/7	3.01×10^{-4}	7.42×10^{-2}

where $\lceil \bar{d} \rceil \equiv \text{round}(\bar{d})$ and diameter D becomes a distance distribution cutoff parameter, $f(d) \ll 1, \forall d > D$ since $f(D)/f(\bar{d}) \sim e^{-100}$. We do not have any singularities that we have to deal with and that are present, for example, in (18). The TZ LS size q is rounded:

$$q = \left\lceil \sqrt{\frac{n}{\log_2 n}} \right\rceil, \quad (22)$$

and all distance distributions are explicitly normalized, e.g. $f(d)$ from (15) is taken to be

$$f(d) = c e^{-\frac{1}{2} \left(\frac{d-\bar{d}}{\sigma} \right)^2}, \quad c \text{ is such that } \sum_{d=1}^D f(d) = 1, \quad (23)$$

and distributions $g(x)$ and $g_1(y)$ are explicitly normalized as well.

For the simulation part, we use our TZ scheme simulator on the graphs produced by the PLRG generator. For a given parameter set, all the data is averaged over 10 random graphs. All average graph sizes n are between 10,000 and 11,000 unless mentioned otherwise.

2.2.1 Distance distribution

We have to stress here that the stretch distribution is a function of the distance distribution and the graph size only. Therefore, all we have to verify for our results having practical value is that both the distance distribution we use for the analysis and the distance distribution in the generated graphs are close to the distance distribution observed in the Internet.

Based on the experiments performed in [85], one can expect that the distance distribution in PLRG-generated graphs should be close to the one in the Internet. We find that it is indeed so. See Fig. 1(a) for details.

Then we proceed as follows. Paying a special attention to the value of the node degree distribution exponent γ equal to 2.1, which is observed in the Internet, we generate series of graphs with γ ranging from 2 to 3, and calculate their distance distributions. We fit these distributions by explicitly normalized Gaussians (23) yielding values of \bar{d} and σ that we use in numerical evaluations of our analytical results. For fitting, we use the standard non-linear least squares method. All fits are very good: the maximum SSE we observe in our fits is 0.003 and the minimum R-square is 0.9905.

The values of \bar{d} and σ in fitted Gaussians are slightly off from the means and standard deviations of distance distributions in generated graphs as depicted in Fig. 1(b). In fact, Fig. 1(b) is a parametric plot of $\sigma(\bar{d})$ with γ being a parameter. We observe almost linear relation between \bar{d} and σ with such parametrization. Note that almost linear relation between the distance c.d.f. center and width parameterized by γ is analytically obtained in [31]. We further discuss this subject in Section 3. In Fig. 1(c,d), we show fitted \bar{d} and σ as functions of γ (cf. the results in [31, 19]).

Average graph sizes for different values of γ are slightly different but dependence of \bar{d} and σ on n (not shown) is negligible compared to their dependence on γ . This is in agreement with [31, 19].

2.2.2 Stretch distribution

We obtain a very close match between the simulations and analysis of the average TZ stretch and stretch distribution. The average stretch as a function of γ is shown in Fig. 2(a). For the Internet-like graphs, $\gamma = 2.1$, the average stretch

Table 3: The average TZ stretch on the $\mathcal{G}_{n,p}$ graphs.

n	p	Avg. degree \bar{k}	(\bar{d}, σ) in graphs	(\bar{d}, σ) in Gaussian fits	\bar{s} (analysis)	\bar{s} (simulations)
10^4	1.3×10^{-3}	13	(3.9, 0.6)	(3.9, 0.5)	1.51	1.60
10^4	5.7×10^{-4}	5.7	(5.5, 0.9)	(5.6, 0.8)	1.37	1.50

we observe in simulations is 1.09 and the average stretch given by (13) with $f(d)$ in (23), with $\bar{d} = 3.4$ and $\sigma = 0.9$, is 1.14.²⁹ Thus, we find that the average stretch is *very low*.

Furthermore, while both the average distance and distance distribution width in power-law graphs do depend on γ (cf. Fig. 1(c,d)), the average stretch *does not*. We delay the discussion of this topic until Section 3.

The stretch distributions obtained both analytically, (14), and in simulations are shown in Fig. 2(b). The sets of significant stretch values (that is, stretch values having noticeable probabilities) match between the analysis and simulations. The top ten stretch values corresponding to virtually 100% of paths are presented in Table 2.

We notice that a majority of paths (up to $\sim 71\%$ according to the simulations) are *shortest*. There are only a very few significant stretch values for the rest of paths. All the significant stretch values are below 2.

The small amount of stretch values with noticeable probabilities is due to the narrow width of the distance distribution. Indeed, in $\sim 86\%$ cases, two random nodes are either 3 or 4 hops away from each other. That is, the probability for x or z to be either 3 or 4 is ~ 0.86 , see Fig. 1(a). In $\sim 82\%$ cases, a random node is just one hop away from its closest landmark, $g_1(1) \sim 0.82$. This explains why stretch-4/3 ($x = 3$, $y = 1$, and $z = 3$) and stretch-5/4 ($x = 4$, $y = 1$, and $z = 4$) paths are most probable among stretch $s > 1$ paths in Table 2.

In Fig. 2(c), the analytical results for the average stretch as a function of the graph size are shown. Note that dependence on n in (13) is only via the LS size q . We present data for the case when \bar{d} and σ are fixed at their values observed in the Internet, and the case when they are allowed to scale as in the DGM model. In both cases, the average stretch slowly *decreases* as the network grows, although this decrease is spread over multiple orders of magnitude of n and the stretch change is confined to a narrow region between 1.3 and 1.1. We also notice that after a certain point, the stretch stops decreasing. Although it becomes very small, it does not reach its minimal value 1.

Finally, in Fig. 2(d), we report the simulation data on the average cluster and LS sizes. We notice that they are well below their bounds. The average cluster size growth similar to the growth of the average distance, cf. Fig. 1(c), is expected.

Recall that the sum of the cluster and LS sizes in the TZ scheme is the number of records in the local routing table. We see that for the Internet-like graphs, $n \sim 10^4$, $\gamma \sim 2.1$, this sum is ~ 52 .

2.2.3 $\mathcal{G}_{n,p}$ graphs

Looking at Figs. 2(a,c), one may be tempted to assume that the average stretch just moderately depends on n and does not depend on either \bar{d} or σ for a wide class of random graphs.

To demonstrate that this is incorrect, we consider the most common class of random graphs, $\mathcal{G}_{n,p}$. We take $n \sim 10^4$ and choose p to match approximately the Internet average distance ($p \sim 1.3 \times 10^{-3}$) and average node degree ($p \sim 5.7 \times 10^{-4}$). The analytical and simulation results for the average stretch in these two cases are presented in Table 3. We find that the average stretch is substantially higher than in the case of random graphs with power-law node degree distributions.

3 The apex

The results in the previous sections suggest that the average stretch depends more strongly on the characteristics of the graph distance distribution (on its first and second moments, in particular) than on the graph size. More specifically, we are taking the distance distribution in a graph to be Gaussian, (23), and, hence, the average TZ stretch \bar{s} in (13) is a function of the average distance \bar{d} and the width of the distance distribution σ in a graph, $\bar{s} \equiv \bar{s}(\bar{d}, \sigma)$. At this point, we wish to explore the analytical structure of $\bar{s}(\bar{d}, \sigma)$ in more detail.

The natural starting point is to fix either \bar{d} or σ in $\bar{s}(\bar{d}, \sigma)$ to their observed values, 3.4 and 0.9 respectively, and to consider two functions, $\bar{s}(\bar{d}, 0.9)$ and $\bar{s}(3.4, \sigma)$, which are shown in Fig. 3(a,b). To our great surprise, we discover that these two functions have *unique minimums* and that the points corresponding to the Internet distance distribution (or, simply, the *Internet points*) are *very close* to them. In other words, one may get an impression that the Internet

²⁹The sources of the small error are in Assumptions 1, 2, and in approximations of Claim 4.

topology has been carefully crafted to result in a distance distribution that would minimize the average TZ stretch. Of course, this can be only an impression and not an explanation since the Internet evolution has had nothing to do with stretch.

The next question we have to ask is if the minimums we observe in Fig. 3(a,b) correspond to a true local minimum of $\bar{s}(\bar{d}, \sigma)$. Our analytical results allow us to collect enough data to draw Fig. 3(c), where the stretch function $\bar{s}(\bar{d}, \sigma)$ is shown for $\bar{d}, \sigma \in [0, 7]$. Note that not all regions of (\bar{d}, σ) correspond to Gaussian-like distance distributions. Indeed, when $\sigma \gtrsim \bar{d}$, distribution $f(d)$ from (23) looks more like an exponential decay since it is cut off from the left by condition $d \geq 1$. Also, when σ is very small and the continuous form of the Gaussian distribution approaches a δ -function, its discrete form (23) goes to either a constant, $f(d) \rightarrow \delta_{d, [\bar{d}]}$, when $\bar{d} \neq (2k + 1)/2$, or to a sum of two constants with equal weights, $f(d) \rightarrow 1/2(\delta_{d,k} + \delta_{d,k+1})$, when $\bar{d} = (2k + 1)/2$. This explains the peculiar peak formation in the $\sigma \sim 0$ area in the picture (see also Appendix C). Networks with such distance distributions are easily constructible (complete networks, stars, various forms of their interconnections, etc.). They all have regular structure, and the exact knowledge of their structure is required for precise stretch calculations. This is why our analytical approach is slightly off in giving the precise answer for stars ($\bar{s} = 1.5$ instead of 1 for $\bar{d} \sim 2, \sigma \sim 0$). Note, however, that for the single case when σ is allowed to be strictly 0, that is, for the complete network case, $\bar{d} = 1, \sigma = 0$, we obtain the correct answer for the average stretch, 2.

The (\bar{d}, σ) -region for distance distributions in realistic networks is, thus, $0 \lesssim \sigma \lesssim \bar{d}$, where we observe a concave area in a form of channel, Fig. 3(c). The area is characterized by particularly low stretch values, which makes us call it the *minimal stretch region* (the MSR). The width and depth of the MSR slowly *increase* as (\bar{d}, σ) grow. In the area of smaller (\bar{d}, σ) , the MSR has a unique critical point, which we call the *MSR apex*. The Internet point is located very close to the apex, which is characterized by the shortest distance between the sets of minimums of \bar{s} —along the \bar{d} - and σ -axes. We may express these sets as two functions, which we denote as $\sigma_{\bar{d}}^*(\bar{d}^*) = \{ (\sigma, \bar{d}) \mid \partial \bar{s} / \partial \bar{d} = 0 \}$ and $\sigma_{\sigma}^*(\bar{d}^*) = \{ (\sigma, \bar{d}) \mid \partial \bar{s} / \partial \sigma = 0 \}$ respectively. We find that $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ *almost touch* each other at the apex. Since the intersection of these two functions would correspond to a stationary point of $\bar{s}(\bar{d}, \sigma)$, we call the apex a *quasi-stationary* point emphasizing that the both derivatives of $\bar{s}(\bar{d}, \sigma)$ are *nearly* zero at this point.

The apex can be more easily observed in Fig. 3(d) showing a projection of Fig. 3(c) on the \bar{d} - σ plane. The solid lines representing the above two sets of minimums forming the MSR, almost touch each other near the apex, and the Internet point is very near their closest segment.

An opportunity to look at the apex from yet another angle is presented in Fig. 3(e) showing a projection of Fig. 3(c) on the \bar{d} - \bar{s} plane. We see that starting from the apex, as \bar{d} increases, the stretch values along $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ become virtually equal and slowly *decrease* as the average distance grows. We also note that $\mathcal{G}_{n,p}$ graphs are far away from the apex and that they have average stretch values that are far from minimal.

We can see now that the apex is indeed a critical or “phase transition” point since it is located at the boundary of the two regions of the average stretch function. The first region, the MSR, is characterized by lowest possible stretch values corresponding to distance distributions observed in real-world graphs. The second region, with substantially higher average stretch values, corresponds to distance distributions in more regular graphs.

To illustrate this point in more detail, we turn our attention back to Fig. 3(d). We see that the two sets of minimums, $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$, are linear in the MSR with sufficiently large $\sigma, \sigma \gtrsim 1$, where the continuous approximation of the distance distribution works particularly well. The two top dashed lines in Fig. 3(d) represent the linear fits of $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ in the area with $\sigma \gtrsim 1$. The exact location of the intersection of these fits is $(\bar{d}^*, \sigma^*) = (3.16, 0.97)$, while the two closest points on the data curves for $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ are $\sigma_{\bar{d}}^*(3.59) = 1.20$ and $\sigma_{\sigma}^*(3.55) = 1.29$. If the linear form of $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ sustained in the area with smaller σ as well, then $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ would intersect at (\bar{d}^*, σ^*) , where we would observe a true stationary point of $\bar{s}(\bar{d}, \sigma)$, which we could then test for the presence of an extremum of the stretch function. This does not happen, however. Instead, as \bar{d} and σ become small, the linear behavior breaks near the apex due to increasingly “more discrete” structure of the distance distribution. See more on this in Appendix C.

In Appendix C, we show that linearity of $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ can be analytically derived from the fact that the distance distribution is taken to be Gaussian. Of course, this does not explain why the *Internet* point is so close either to the MSR or to its apex.

The linear form of $\sigma_{\bar{d}}^*(\bar{d}^*)$ and $\sigma_{\sigma}^*(\bar{d}^*)$ sheds some light on a closely related issue of why the average stretch is virtually independent of γ . In Fig. 3(d), the shaded area represents a set of (\bar{d}, σ) , for which the average stretch is approximately the same as for the Internet, $S_I = \{ (\bar{d}, \sigma) \mid \bar{s}(\bar{d}, \sigma) \sim \bar{s}(3.4, 0.9) \}$. In other words, it is a projection of the cyan area in Fig. 3(c) on the \bar{d} - σ plane. We see that in the MSR, the S_I boundaries are almost parallel straight

lines. Therefore, if the average stretch is to be independent of γ , which is observed in Section 2.2.2, then the points representing distance distributions in power-law graphs, $(\bar{d}_\gamma, \sigma_\gamma)$, from Fig. 1(b) should lie along the S_I boundaries, and this is what indeed happens. Yet again, the linear relation between \bar{d}_γ and σ_γ in the power-law graphs, and the fact that this relation is just as required for the average TZ stretch being virtually independent of γ , come from two seemingly disjoint domains.

To finish the list of various “coincidences,” we construct a linear fit of $(\bar{d}_\gamma, \sigma_\gamma)$ (the bottom-most dashed line in Fig. 3(d)). The Internet point, $\gamma = 2.1$, lies on this line. Our numeric analysis shows that the Internet value of $\gamma = 2.1$ is a unique value of γ minimizing the distance between the linear fit of $(\bar{d}_\gamma, \sigma_\gamma)$ and (\bar{d}^*, σ^*) , which is the intersection of the linear fits of $\sigma_d^*(\bar{d}^*)$ and $\sigma_\sigma^*(\bar{d}^*)$. In other words, the Internet distance distribution is the point that is *closest* to the MSR apex, compared to distance distributions in all other scale-free graphs with power-law node degree distributions.

4 Conclusions

Of course, the TZ scheme, reducing, in principle, the routing table size to about 50 entries for 10^4 -node scale-free networks, and making routing decision running time constant, cannot pretend to be a realistic Internet interdomain routing scheme. First, it is static. Second, addressing in interdomain routing is based on IP addresses rather than on interdomain graph node labels, that is, AS numbers.³⁰ Third, it assumes availability of the global topology view.

Most importantly, in the context of our work, the scheme is not a stretch-1 scheme. Indeed, interdomain routing in the Internet is essentially shortest path routing.³¹ A routing scheme that would prevent a pair of ASs from utilizing a peering link between them is not realistic, of course. Thus, any stretch $s > 1$ routing scheme applied to the Internet would involve augmentation, in one form or another, of the routing information provided by the scheme with the shortest path routing information for non-shortest paths. This explains why we are concerned with the average stretch produced by a scheme.

Our principal finding that the average TZ stretch on the Internet graph is reasonably low opens a well-defined path for the future work in the area of applying relevant theoretical results obtained for routing to realistic scale-free networks (see the next section). If the average stretch of even the “exceptional” TZ scheme turned out to be relatively high, the scheme would be inapplicable *in principle* (not just *in practice*), which would essentially close the above path, demonstrate impossibility to construct efficient and scalable routing for the Internet, and call for searching one somewhere beyond the traditional graph-theoretical approach.

As we mention in the introduction, our finding that the Internet distance distribution is in a close neighborhood of the MSR apex cannot be explained in the present idea set since the Internet, as we know it today, has nothing to do with stretch. While we lack sufficient information to show cause for this effect, we do believe it strongly suggests the analytical structure of the average stretch function may be an indirect (or even direct) indicator of some yet-to-be discovered processes that have influenced the Internet’s topological evolution. In other words, a rigorous explanation of this phenomenon would probably require much deeper understanding of the Internet evolution principles (that are far from being even known if we accept the critique of the BA model and alike) and demonstration of a link between them and the TZ scheme.

We believe that an explanation of this effect will most probably have the following pattern. The Internet evolution principles turn out to be such that they minimize X , where X is some known or yet unknown characteristic of a network. At the same time, the distance to the MSR apex turns out to be a monotonically increasing function of X . There are reasons to believe that the distance to the apex is not something random. Indeed, as we have seen, the apex is a *unique* critical point of the average TZ stretch function, and, at the same time, the TZ scheme is also “exceptional” in a sense that it delivers a nearly optimal first possibility to deviate from incompressible shortest path routing. The outlined pattern would create a necessary link between the two seemingly disjoint domains.

5 Future work

The list of immediate practical next questions that remain open include:

- What is the average stretch produced by the *Cowen* scheme on scale-free graphs?

³⁰Although there are some proposals, *Atomized Routing*, [11, 12, 91], and *ISLAY*, [54], suggesting to “fix” this. If this is “fixed,” we should be ready to face the problem of accelerated rates of growth of the total number of ASs.

³¹Shortest path routing in the Internet is perturbed by various administrative constraints called *policies*, the routing protocol, BGP, being a policy routing tool. For measurements of stretch produced by policy routing, see [86].

Obtaining analytical results for the Cowen scheme is a harder problem. However, one can expect that the average stretch would be lower. Indeed, it is easy to see that making high-betweenness nodes belong to the LS should decrease the average stretch. The Cowen LS construction algorithm uses the greedy set cover algorithm by Lovász preferring high degree nodes, [64], but as observed in [90], betweenness is linearly correlated with node degrees in the Internet. The expectation of the Cowen scheme producing an average stretch that is lower than in the TZ case is in a good agreement with the general stretch-space trade-off since the Cowen scheme has a higher local memory space upper bound.

- Do any routing schemes based on the *additive* stretch factor deliver lower average stretch on scale-free graphs? The multiplicative stretch factor may be too coarse for short distances that prevail in scale-free graphs.

- What is the memory space lower bound for *shortest path* routing on scale-free graphs?

We know that generic shortest path routing is incompressible. However, the situation is better for almost all graphs. We do not know any bounds for scale-free graphs. Answering this question has very important practical implications since (policy-constrained) shortest path routing seems to be a requirement for Internet interdomain routing.

- More specifically, can any upper bounds be obtained for the total number of stretch $s > 1$ paths produced by existing $s > 1$ routing schemes on scale-free graphs?

If these upper bounds are found to be as low as the total space upper bounds, then the original routing information can be augmented with the $s = 1$ information for $s > 1$ paths without increasing the space upper bounds.

- What bounds can be obtained for *dynamic* routing on scale-free graphs?

Of course, no realistic Internet routing can neglect the adaptation cost considerations. While of critical practical importance, obtaining various bounds for dynamic routing on scale-free graphs seems to be the hardest problem among those we list here. Furthermore, one can hardly expect such bounds to be low, taking into consideration rather pessimistic lower bounds obtained for dynamic routing on generic networks (cf. Section 1.1.1). In addition, it is clear that the TZ scheme cannot be easily modified to perform well in the dynamic case since the scheme labels nodes with topology-sensitive information. In other words, the scheme is not name-independent. As soon as topology changes, nodes need to be relabelled. Significant progress in construction of name-independent static low-stretch routing schemes has been recently made by Arias, Cowen, et al. in [7].

On the theoretical side, which may turn out to have practical implications as well, the explanation of Internet distance distribution proximity to the MSR apex appears to be a very interesting problem. It would be much easier to solve, of course, if the integral in (18) could be evaluated analytically. What might be an alternative set of assumptions that would make an expression analogous to (18) analytically solvable? Can similar results be obtained for other forms of distance distributions, and, yet more importantly, for other routing schemes (the Cowen scheme, or the stretch-5 routing scheme with the $\tilde{O}(n^{1/2})$ local memory space upper bound obtained in [33], for example)? A technical issue with *experimental* (vs. analytical) studies of the MSR is that we do not know an efficient algorithm to generate graphs with a given distance distribution.

In any case, the explanation of Internet’s proximity to the apex is of great theoretical interest, as the fundamental laws governing the Internet evolution remain unclear. Therefore, on the practical side, a proper explanation of this effect may help us, for example, in our intent to move, [17], from purely *descriptive* Internet evolution models to more *explanatory* ones, in the terminology of the program outlined by the authors of [94].

The Internet started as a small research network designed and fully controlled by a group of few enthusiasts ([62]). Today, after a series of “phase transitions,” it has evolved to a huge network interconnecting tens of thousands of independent and even adversarial networks without a single point of external control. This makes the Internet a “self-governing,” “self-evolving” complex system, a research subject of areas of physics (statistical mechanics, in particular) studying evolution of such systems in general. Construction of realistic, efficient, and scalable routing for this “new” Internet is an interesting and challenging task lying ahead of us.

References

- [1] Y. Afek, E. Gafni, and M. Ricklin. Upper and lower bounds for routing schemes in dynamic networks (abstract). In *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 370–375. IEEE, 1989.

- [2] W. Aiello, F. Chung, and L. Lu. A random graph model for massive graphs. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 171–180. ACM Press, 2000.
- [3] D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal on Computing*, 28(4):1167–1181, 1999.
- [4] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.
- [5] R. Albert, H. Jeong, and A.-L. Barabási. The diameter of the World Wide Web. *Nature*, 401:130–131, 1999.
- [6] L. A. N. Amaral, A. Scala, M. Barthélémy, and H. E. Stanley. Classes of small-world networks. *PNAS*, 97:11149–11152, 2000.
- [7] M. Arias, L. Cowen, K. A. Laing, R. Rajaraman, and O. Taka. Compact routing with name independence. In *Proceedings of the 15th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 184–192. ACM, 2003.
- [8] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [9] BGP Table Data, <http://bgp.potaroo.net/>.
- [10] B. Bollobás. *Random Graphs*. Academic Press, New York, 1985.
- [11] A. Broido and k claffy. Analysis of RouteViews BGP data: Policy atoms. In *Proceedings of Network Resource Data Management Workshop (NRDM)*, 2001.
- [12] A. Broido and k claffy. Complexity of global routing policies. In *Cooperative Association for Internet Data Analysis (CAIDA)*, 2001. <http://www.caida.org/projects/routing/atoms/>.
- [13] A. Broido and k claffy. Internet topology: Connectivity of IP graphs. In *SPIE International Symposium on Convergence of IT and Communication*, 2001.
- [14] A. Broido, E. Nemeth, and k claffy. Internet expansion, refinement, and churn. *European Transactions on Telecommunications*, 13(1):33–51, 2002.
- [15] H. Buhrman, J.-H. Hoepman, and P. Vitány. Space-efficient routing tables for almost all networks and the incompressibility method. *SIAM Journal on Computing*, 28(4):1414–1432, 1999.
- [16] I. Castineyra, N. Chiappa, and M. Steenstrup. *The Nimrod Routing Architecture*. IETF, RFC 1992, 1996.
- [17] H. Chang, S. Jamin, and W. Willinger. Internet connectivity at the AS level: An optimization driven modeling approach. In *Proceedings of MoMeTools*, 2003.
- [18] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. The origin of power laws in Internet topologies revisited. In *IEEE INFOCOM*, 2002.
- [19] F. Chung and L. Lu. The average distance in a random graph with given expected degrees. *Internet Mathematics*, 1(1):91–114, 2003.
- [20] R. Cohen, D. Dolev, S. Havlin, T. Kalisky, O. Mokryn, and Y. Shavitt. On the tomography of networks and multicast trees. Technical Report TR49, Hebrew University, 2002.
- [21] R. Cohen and S. Havlin. Scale-free networks are ultrasmall. *Physical Review Letters*, 90(05):058701, 2003.
- [22] L. Cowen. Compact routing with minimum stretch. *Journal of Algorithms*, 38(1):170–183, 2001.
- [23] S. Dolev, E. Kranakis, D. Krizanc, and D. Peleg. Bubbles: Adaptive routing scheme for high-speed dynamic networks. *SIAM Journal on Computing*, 29(3):804–833, 1999.
- [24] D. Dor, S. Halperin, and U. Zwick. All pairs almost shortest paths. *SIAM Journal on Computing*, 29:1740–1759, 2000.
- [25] A. Doria, E. Davies, and F. Kastenzholz. *Requirements for Inter-Domain Routing*. IRTF, Internet Draft, 2003.

- [26] S. N. Dorogovtsev. Private communications.
- [27] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes. Pseudofractal scale-free web. *Physical Review E*, 65(06):066122, 2002.
- [28] S. N. Dorogovtsev and J. F. F. Mendes. Language as an evolving Word Web. *Proceedings of the Royal Society of London B*, 268:2603–2606, 2001.
- [29] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks. *Advances in Physics*, 51:1079–1187, 2002.
- [30] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press, Oxford, 2003.
- [31] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin. Metric structure of random networks. *Nuclear Physics B*, 653(3):307–422, 2003.
- [32] Y. Dourisboure and C. Gavoille. Improved compact routing scheme for chordal graphs. In *Proceedings of the 16th International Conference on Distributed Computing (DISC)*, volume 2508 of *Lecture Notes in Computer Science*, pages 252–264. Springer, 2002.
- [33] T. Eilam, C. Gavoille, and D. Peleg. Compact routing schemes with low stretch factor. *Journal of Algorithms*, 46:97–114, 2003.
- [34] M. Elkin and D. Peleg. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 173–182. ACM, 2001.
- [35] P. Erdős and A. Rényi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- [36] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the Internet. In *Proceeding of the 29th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2380 of *Lecture Notes in Computer Science*, pages 110–122. Springer, 2002.
- [37] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of the ACM SIGCOMM*, pages 251–262, 1999.
- [38] F. Fich. End-to-end communication. In *Proceedings of the 2nd International Conference on Principles of Distributed Systems (OPODIS)*, pages 37–43. Hermes, 1998.
- [39] M. Flammini, J. van Leeuwen, and A. Marchetti-Spaccamela. The complexity of interval routing on random graphs. *The Computer Journal*, 41(1):16–25, 1998.
- [40] P. Fraigniaud and C. Gavoille. Lower bounds for oblivious single-message end-to-end communication. Research Report RR-1289-03, LaBRI, University of Bordeaux, 2003.
- [41] A. Fronczak, P. Fronczak, and J. A. Hołyst. Average path length in random networks. cond-mat/0212230.
- [42] C. Gavoille. A survey on interval routing. *Theoretical Computer Science*, 245(2):217–253, 2000.
- [43] C. Gavoille. Routing in distributed networks: Overview and open problems. *ACM SIGACT News - Distributed Computing Column*, 32(1):36–52, 2001.
- [44] C. Gavoille and M. Genegler. Space-efficiency for routing schemes of stretch factor three. *Journal of Parallel and Distributed Computing*, 61(5):679–687, 2001.
- [45] C. Gavoille and M. Néhez. Interval routing in reliability networks. In *Proceeding of the 9th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*. Carleton University Press, 2003. To appear.
- [46] C. Gavoille and D. Peleg. The compactness of interval routing for almost all graphs. In *Proceedings of the 12th International Symposium on Distributed Computing (DISC)*, volume 1499 of *Lecture Notes in Computer Science*, pages 161–174. Springer, 1998.
- [47] C. Gavoille and S. Pérennès. Memory requirement for routing in distributed networks. In *Proceedings of the 15th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 125–133. ACM, 1996.

- [48] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, San Diego, sixth edition, 2000.
- [49] J. Guare. *Six Degrees of Separation: A Play*. Vintage Books, New York, 1990.
- [50] G. Huston. Analyzing the Internet’s BGP routing table. *The Internet Protocol Journal*, 4(1), 2001.
- [51] G. Huston. *Commentary on Inter-Domain Routing in the Internet*. IETF, RFC 3221, 2001.
- [52] G. Huston. Scaling inter-domain routing—a view forward. *The Internet Protocol Journal*, 4(4), 2001.
- [53] H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A.-L. Barabási. The large-scale organization of metabolic networks. *Nature*, 407:651–654, 2000.
- [54] F. Kastenholz. *ISLAY: A New Routing and Addressing Architecture*. IRTF, Internet Draft, 2002.
- [55] L. Kleinrock and F. Kamoun. Hierarchical routing for large networks: Performance evaluation and optimization. *Computer Networks*, 1:155–174, 1977.
- [56] L. Kleinrock and F. Kamoun. Optimal clustering structures for hierarchical topological design of large computer networks. *Networks*, 10:221–248, 1980.
- [57] D. Krioukov and D. Volchenkov. An economy-based model for the Internet evolution, 2003. To appear.
- [58] D. Krizanc, F.L. Luccio, and R. Raman. Compact routing schemes for dynamic ring networks. Technical Report 2002/26, Department of Mathematics and Computer Science, University of Leicester, 2002.
- [59] A. Krzywicki. Defining statistical ensembles of random networks. In *Workshop on Discrete Random Geometries and Quantum Gravity*, Utrecht, 2001. cond-mat/0110574.
- [60] E. Kushilevitz, R. Ostrovsky, and A. Rosén. Log-space polynomial end-to-end communication. *SIAM Journal on Computing*, 27(6):1531–1549, 1998.
- [61] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proceedings of IEEE INFOCOM*, 2003.
- [62] B. Leiner, V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts, and S. Wolff. A brief history of the Internet. *Communications of the ACM*, 40(2):102–108, 1997.
- [63] F. Liljeros, C. R. Edling, L. A. N. Amaral, H. E. Stanley, and Y. Aberg. The web of human sexual contacts. *Nature*, 411:907–908, 2001.
- [64] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
- [65] B. Mandelbrot. An informational theory of the statistical structure of languages. In *Communication Theory*, pages 486–502. Butterworth, 1953.
- [66] S. Milgram. The small world problem. *Psychology Today*, 61, 1967.
- [67] M. Molloy and B. Reed. A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, 6:161–179, 1995.
- [68] M. Molloy and B. Reed. The size of the giant component of a random graph with a given degree sequence. *Combinatorics, Probability and Computing*, 7:295–305, 1998.
- [69] J. M. Montoya and R. V. Sole. Small world patterns in food webs. Technical Report 00-10-059, Santa Fe Institute, 2000.
- [70] M. E. J. Newman. The structure of scientific collaboration networks. *PNAS*, 98:404–409, 2001.
- [71] M. E. J. Newman. Random graphs as models of networks. In *Handbook of Graphs and Networks*, Berlin, 2003. Wiley-VCH.
- [72] M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64(02):026118, 2001.

- [73] M. E. J. Newman, D. J. Watts, and S. H. Strogatz. Random graph models of social networks. *PNAS*, 99:2566–2572, 2002.
- [74] W. B. Norton. Internet Service Providers and peering. Equinix white paper.
- [75] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA, 2000.
- [76] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, 36:510–530, 1989.
- [77] R. Perlman. Hierarchical networks and subnetwork partition problem. *Computer Networks and ISDN Systems*, 9:297–303, 1985.
- [78] E. Ravasz and A.-L. Barabási. Hierarchical organization in complex networks. *Physical Review E*, 67(02):026112, 2003.
- [79] E. Ravasz, A. L. Somera, D. A. Mongru, Z. N. Oltvai, and A.-L. Barabási. Hierarchical organization of modularity in metabolic networks. *Science*, 297:1551–1555, 2002.
- [80] S. Ross. *Introduction to Probability Models*. Academic Press, San Diego, seventh edition, 2000.
- [81] Self-Organized Networks at Notre Dame, <http://www.nd.edu/~networks/>.
- [82] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos. Power-laws and the AS-level Internet topology. In *ACM/IEEE Transactions on Networking*, 2003. To appear.
- [83] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proceeding of IEEE INFOCOM*, 2002.
- [84] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, W. Willinger, and S. Shenker. Does AS size determine AS degree? *ACM Computer Communication Review*, October 2001.
- [85] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: Degree-based vs. structural. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 147–159. ACM Press, 2002.
- [86] H. Tangmunarunkit, R. Govindan, and S. Shenker. Internet path inflation due to policy routing. In *SPIE ITCOM*, 2001.
- [87] M. Thorup and U. Zwick. Approximate distance oracles. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 183–192. ACM, 2001.
- [88] M. Thorup and U. Zwick. Compact routing schemes. In *Proceedings of the 13th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 1–10. ACM, 2001.
- [89] A. Vázquez, R. Pastor-Satorras, and A. Vespignani. Internet topology at the router and Autonomous System level. `cond-mat/0206084`.
- [90] A. Vázquez, R. Pastor-Satorras, and A. Vespignani. Large-scale topological and dynamical properties of the Internet. *Physical Review E*, 65(06):066130, 2002. `cond-mat/0112400`.
- [91] P. Verkaik, A. Broido, and k claffy. Atomized routing. In *Fourth International Workshop on the Design of Reliable Communication Networks (DRCN)*, 2003. Submitted.
- [92] D. J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, Princeton, NJ, 1999.
- [93] D. J. Watts and S. H. Strogatz. Collective dynamics of “small-world” networks. *Nature*, 393:440–442, 1998.
- [94] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker. Scaling phenomena in the Internet: Critically examining criticality. *PNAS*, 99(Suppl. 1):2573–2580, 2002.

Appendices

A The Kleinrock-Kamoun stretch on the Internet

In this appendix, we calculate a rough estimate of the stretch factor of the Kleinrock-Kamoun (KK) hierarchical routing scheme, [55], applied to the observed Internet interdomain topology. We find that the stretch is very high, which is consistent with the observation made in [55] that the approach used there works reasonably well only for sparsely connected networks. The scale-free networks, on the contrary, are extremely densely connected.

Recall that [55] assumes existence of a hierarchical partitioning of a network of size n into m levels of clusters. Each k -level cluster consists of $n^{1/m}$ ($k-1$)-level clusters, $k = 1 \dots m$, 0-level clusters being nodes. The optimal clustering is achieved when $m \sim \log n$. There are few other fairly strong assumptions about the properties of required partitioning. Neither algorithm for its construction nor proof of its existence are delivered, but if it does exist then the stretch factor is shown to be

$$s = 1 + \frac{1}{\bar{d}} \sum_{k=1}^{m-1} \left[1 - \frac{n^{\frac{k}{m}} - 1}{n - 1} \right] d_k, \quad (24)$$

where \bar{d} is the network average distance and d_k is the diameter of a k -level cluster.

It is further assumed in [55] that both the network diameter and average distance are power-law functions of the network size. This is certainly not true for scale-free networks with power-law *node degree* distributions. For the very recent results on the average distance in such networks see the references mentioned in Section 1.1.2. In the numerical evaluations in this appendix, we use the value of $\bar{d} \sim 3.6$ observed in the Internet.

Much fewer results are available for the network diameter. However, in [19], it is shown that the diameter of networks with power-law node degree distribution with exponent γ lying between 2 and 3 scales almost surely as $\Theta(\log n)$. For the Internet, $\gamma \sim 2.1$, and since the Internet size $n \sim 1.5 \times 10^4$ is relatively large, we may write the Internet diameter D as $D \sim c \log n$ with some multiplicative coefficient c . The observed value of D , $D \sim 13$ ([9]), defines c then.

The size of a k -level cluster is obviously $n^{k/m}$ but nothing rigorous can be said about its degree distribution since there is no procedure for its construction. Thus, it is natural to assume that its degree distribution is also power-law with $2 < \gamma < 3$, which gives an estimate of the k -level cluster diameter as $d_k \sim c \log n^{k/m} \sim Dk/m$. Substituting this in (24) and performing summation gives

$$s \sim 1 + \frac{D}{2\bar{d}} \left[m \frac{n}{n-1} - \frac{n(n^{\frac{2}{m}} - 1)}{(n-1)(n^{\frac{1}{m}} - 1)^2} + \frac{2}{m} \frac{n^{\frac{1}{m}}}{(n^{\frac{1}{m}} - 1)^2} \right]. \quad (25)$$

Using the numerical values for n , \bar{d} , D , and optimal $m = 10$, we can see that the KK stretch factor on the Internet interdomain topology is

$$s \sim 15. \quad (26)$$

Note that a 15-time path length increase in the Internet would lead to AS path lengths of ~ 55 and IP hop path lengths of ~ 150 .

Stretch factors for smaller, non-optimal values of m are shown in Fig. 4. The fact that the stretch grows almost linearly with the number of hierarchical levels follows directly from (25), which, for large n , can be rewritten as

$$s \sim 1 + \frac{D}{2\bar{d}}(m - 1). \quad (27)$$

Using $D \sim \log n$, $\bar{d} \sim \log \log n$ ([19]), and optimal $m \sim \log n$, we obtain the following estimate of the stretch factor as a function of the network size:

$$s \sim \frac{\log^2 n}{\log \log n}. \quad (28)$$

B Proofs

In this appendix, we prove various statements made in Section 2.1.

B.1 Claims 1 and 2

A rigorous proof would involve the same type of argument as used to show that the hypergeometric distribution converges to the binomial distribution for small selections from large sets. We provide a close outline of the proof.

Suppose we have a set of n objects of two types: n_1 objects of type 1 and n_2 objects of type 2, $n_1 + n_2 = n$ and $\vartheta = n_1/n$. Recall that the hypergeometric distribution,

$$p_h(x) = \frac{\binom{n_1}{x} \binom{n_2}{m-x}}{\binom{n}{m}}, \quad (29)$$

gives the probability that a random sample of m objects from the set of all objects contains x objects of type 1. If sampling is with replacement—that is, as soon as one object is selected, it is immediately returned back to the set before the next object is selected,—then the probability that x type-1 objects have been selected after m one-object selections is given by the binomial distribution,

$$p_b(x) = \binom{m}{x} \vartheta^x (1 - \vartheta)^{m-x}. \quad (30)$$

One can easily see that probabilities for small selections with and without replacement converge in the limit of large n ,

$$p_h(x) \xrightarrow[m/n \rightarrow 0]{n \rightarrow \infty} p_b(x). \quad (31)$$

Suppose now that $f(d)$ and $F(d)$ are distance p.d.f. and c.d.f. of some graph of size n , diameter D , and that the graph LS L is of size q . The problem of finding the probability $g_i(d)$ that the i 'th closest landmark is at distance d is equivalent to the following object selection problem. The total number of objects to select from is n , the number of types of objects is D , and the number of objects of type d , $d = 0 \dots D$, is $f(d)n$. In addition, q random objects are marked as landmarks. Since landmarks are randomly marked, $g_i(d)$ is the probability that a randomly selected object is of type d , which is $f(d)$ by definition, times the probability that $i - 1$ random objects are of (possibly different) types $d_- \leq d$ (corresponding to the closer landmarks), times the probability that $q - i$ random objects are of (possibly different) types $d_+ \geq d$ (corresponding to the farther landmarks). That is, denoting the latter two probabilities by $p_-(d)$ and $p_+(d)$ respectively, we can write that

$$g_i(d) = c_i p_-(d) f(d) p_+(d) \quad (32)$$

with some normalization coefficient c_i .

The probability than one random object is of type d_- (d_+) is $F(d)$ ($1 - F(d)$) by definition, but obtaining the exact answer for $p_-(d)$ ($p_+(d)$) involves some combinatorics resulting in a generalization of the hypergeometric distribution (29), which is left as an exercise for a rigorous reader. However, since $q/n \xrightarrow[n \rightarrow \infty]{} 0$, sampling without replacement can be approximated by sampling with replacement, cf. (31), which leads to the following approximations for $p_-(d)$ and $p_+(d)$:

$$p_-(d) \rightarrow F(d)^{i-1}, \quad (33)$$

$$p_+(d) \rightarrow (1 - F(d))^{q-i}. \quad (34)$$

The analogy to the order statistics becomes straightforward now.

The normalization coefficient c_i in (32) is defined by the normalization condition

$$1 = \sum_{d=0}^D g_i(d) = c_i \sum_{d=0}^D F(d)^{i-1} f(d) (1 - F(d))^{q-i}, \quad (35)$$

which can be expressed as the following Lebesgue-Stieltjes integral:

$$1 = c_i \int_0^D F(d)^{i-1} (1 - F(d))^{q-i} dF(d) = c_i \int_0^1 F^{i-1} (1 - F)^{q-i} dF. \quad (36)$$

In the continuous limit, evaluation of the corresponding Riemann integral gives

$$1 = c_i \frac{\Gamma(i) \Gamma(q - i + 1)}{\Gamma(q + 1)} = c_i \left[i \binom{q}{i} \right]^{-1}. \quad (37)$$

B.2 Claim 3

If we denote by $\tilde{G}_1(d)$ the c.c.d.f. for the distance to the closest landmark,

$$\tilde{G}_1(d) = \sum_{\tilde{d}=d}^D g_1(\tilde{d}), \quad (38)$$

then according to the cluster definition (1), the average cluster size is

$$|\overline{C}| = \sum_{d=0}^D n f(d) \tilde{G}_1(d). \quad (39)$$

Using the explicit expression for g_1 in (3), we can calculate an upper bound for \tilde{G}_1 by representing the sum in (38) as the following Lebesgue-Stieltjes integral:

$$\tilde{G}_1(d) = q \int_d^D (1 - F(\tilde{d}))^{q-1} dF(\tilde{d}) = q \int_{F(d)}^1 (1 - F)^{q-1} dF. \quad (40)$$

Evaluation of the corresponding Riemann integral gives

$$\tilde{G}_1(d) \leq (1 - F(d))^q, \quad (41)$$

where equality is attained in the continuous limit. Note that the obtained upper bound, $(1 - F(d))^q$, is equivalent to the c.c.d.f. for the geometric distribution, $p_g(x) = \vartheta(1 - \vartheta)^x$, with the success probability $\vartheta = F(d)$ and number of trials $x = q - 1$. The reason for that becomes transparent after formulating the problem in the object selection framework considered in Appendix B.1.

Substituting (41) in (39) completes the proof:

$$|\overline{C}| = n \int_0^D \tilde{G}_1(d) dF(d) \leq n \int_0^1 (1 - F)^q dF \leq \frac{n}{q+1}. \quad (42)$$

Note that clusters are originally defined as objects “inverse” to balls $B(v) = \{ b \in V \mid \delta(b, v) < \delta(v, L(v)) \}$. They are “inverse” in the following sense: $v \in B(w) \Leftrightarrow w \in C(v)$. If we denote the c.d.f. of $g_1(d)$ by $G_1(d)$, then the average ball size is

$$|\overline{B}| = \sum_{n=0}^D n g_1(d) F(d) = n \int_0^D F(d) dG_1(d). \quad (43)$$

Since (39) can be written as

$$|\overline{C}| = n \left(1 - \int_0^D G_1(d) dF(d) \right), \quad (44)$$

the known fact that the average ball and cluster sizes are equal is just a consequence of integration by parts of the following Lebesgue-Stieltjes integral:

$$1 = F(d)G_1(d) \Big|_0^D = \int_0^D F(d) dG_1(d) + \int_0^D G_1(d) dF(d). \quad (45)$$

B.3 Claim 4

Note that the first case in (10), $z < y$, accounts exactly for stretch-1 paths from w to destinations $v \in C(w)$, (see definition (1)).

The second case, $z < x$, accounts for the degenerate shortcut on paths from w to $L(v)$ that go through v . On such paths, $x = z + y$, the length of the shortcut portion of the total path from w to v is zero, and the stretch is

1. It is easy to see that, on small-world graphs, paths from w to $L(v)$ that do not go through v but that still have shortcuts are rare and hard to account for without knowing the graph topology.

Indeed, consider triangle $\triangle WVL$ in Fig. 5, where W represents source node w , V is destination v , and L is $L(v)$. Shortcutting occurs when destination V is “in front of” its landmark L , which is represented by placing V to the left of the vertical line. It is clear that on small-world graphs, characterized by low average distances and narrow distance distributions, the possibility for V to *not* lie on path WL is minuscule.

To see this, suppose we still want to account for cases when V does not belong to WL by approximating such cases within the 2-dimensional Euclidean space. The shortcut path is then WUV in Fig. 5, with U representing the first node u on the path from w to $L(v)$ such that $v \in C(u)$. Its position on WL is defined by $|UV| = |VL| = y$, and the length of segment WU is denoted by x^* . The stretch is then given by $s^* = (x^* + y)/z$, and solving $\triangle WUV$, we find that $x^* = (z^2 - y^2)/x$. Our numerical experiments show that it does not matter if, with condition $z < x$, we use $s^* = (x^* + y)/z$ or $s^* = 1$: the two stretch distributions are virtually identical and the average stretch difference is only in the fourth digit. This means that $(x^* + y)/z$ is virtually always 1 as soon as $z < x$. From $(x^* + y)/z = [(z^2 - y^2)/x + y]/z = 1$ follows $x = z + y$, which means that V lies on WL .

Note, however, that condition $z < x$ is not the exact condition for the shortcut presence in the Euclidean space. As mentioned above, the exact condition is that V is to the left of the vertical line, which implies $z^2 < x^2 + y^2$. In our experiments, we observe that using this condition instead of $z < x$ leads to the stretch distribution drastically distinct from the one observed in simulations. This is because any estimates based on approximation of the finite metric space in a graph by the Euclidean metric space are not applicable to small-world graphs. Such estimates are more applicable to grid-like graphs with wide distance distributions and average distances growing as power-law functions of the graph size.

B.4 Claim 5

Suppose the distance matrix \mathcal{D} is given. Let us denote by $\mathcal{D}(\alpha)$ the set of elements in \mathcal{D} equal to α (α -elements),

$$\mathcal{D}(\alpha) = \{ \mathcal{D}_{ij}; i, j = 1 \dots n \mid \mathcal{D}_{ij} = \alpha \}. \quad (46)$$

Note that the distance p.d.f. is equal to the distribution of values of elements in \mathcal{D} ,

$$f(\alpha) = \frac{|\mathcal{D}(\alpha)|}{|\mathcal{D}|}. \quad (47)$$

We can also define the set of elements satisfying the triangle inequality,

$$\mathcal{D}(\beta, \gamma) = \{ \mathcal{D}_{ij}; i, j = 1 \dots n \mid |\beta - \gamma| \leq \mathcal{D}_{ij} \leq \beta + \gamma \} = \bigcup_{\alpha=|\beta-\gamma|}^{\beta+\gamma} \mathcal{D}(\alpha). \quad (48)$$

The distribution of α -elements in $\mathcal{D}(\beta, \gamma)$ is

$$p(\alpha|\beta, \gamma) = \frac{|\mathcal{D}(\alpha)|}{|\mathcal{D}(\beta, \gamma)|} = \frac{f(\alpha)}{\sum_{\delta=|\beta-\gamma|}^{\beta+\gamma} f(\delta)} \quad (49)$$

Suppose that a group of k nodes is randomly selected in the graph, and that their indices are $\mathbf{i}_k = \{i_1, i_2, \dots, i_k\}$. The distribution of distances between them is also the distribution of values of elements in a $k \times k$ submatrix $\mathcal{D}_{\mathbf{i}_k}$ obtained from \mathcal{D} by intersecting the \mathbf{i}_k rows and columns. Since the selection is random, this distribution is also $f(\alpha)$. If *two* groups of random nodes, \mathbf{i}_{k_1} and \mathbf{j}_{k_2} , are selected and we are looking for the distribution of distances between pairs of nodes belonging to the *different* groups, we construct a set $\mathcal{D}_{\mathbf{i}_{k_1}, \mathbf{j}_{k_2}}$ by intersecting the \mathbf{i}_{k_1} rows of \mathcal{D} with the \mathbf{j}_{k_2} columns and *vice versa*. This subset of \mathcal{D} is no longer a submatrix of \mathcal{D} , but the distribution of α -elements in it is still $f(\alpha)$.

Our problem is to find the conditional distribution $t(z|x, y)$ for distance z between a group of $g(x)n$ (on average) random nodes at distance x from some random node (the landmark closest to the destination) and another group of $g_1(y)n$ (on average) nodes at distance y from the same node. The groups may overlap when $x = y$. They define the subset $\mathcal{D}_{\mathbf{i}_{g(x)n}, \mathbf{j}_{g_1(y)n}}$ as in the previous paragraph. The distribution we are looking for, $t(z|x, y)$, is the distribution of values of elements in this subset. Since there are no correlations in the distance matrix, the only difference between this subset and truly random subsets $\mathcal{D}_{\mathbf{i}_{k_1}, \mathbf{j}_{k_2}}$ from the previous paragraph is due to the triangle inequality caused by the fact that the two groups of nodes are neighbors of the same node. This means that $\mathcal{D}_{\mathbf{i}_{g(x)n}, \mathbf{j}_{g_1(y)n}}$ is also a

subset of $\mathcal{D}(x, y)$ defined in (48), and, since $\mathcal{D}_{\mathbf{1}_{g(x)n} \mathbf{1}_{g_1(y)n}}$ is random in other respects, the distribution of values of its elements is the same as in $\mathcal{D}(x, y)$. In other words, $t(z|x, y) = p(z|x, y)$ in (49).

Noticing that, by the formula for the conditional probability, $t(x, y, z) = t(z|x, y)g(x)g_1(y)$ completes the proof.

C The MSR analysis

In this appendix we demonstrate that linearity of functions $\sigma_{\bar{d}, \sigma}^*(\bar{d}^*)$ in the MSR follows from the Gaussian form of the distance distribution.

We consider the continuous case, for which $\bar{s}(\bar{d}, \sigma)$ is given by (18). The necessary condition for a local minimum of \bar{s} in the \bar{d} - or σ -direction is $\partial \bar{s} / \partial \bar{d} = 0$ or $\partial \bar{s} / \partial \sigma = 0$ respectively. After some algebra using $\operatorname{erf}'(\alpha) = 2e^{-\alpha^2} / \sqrt{\pi}$, we obtain, with x, y , and z having the same semantics as in Section 2.1,

$$\begin{aligned} \frac{\partial \bar{s}}{\partial \bar{d}} = 0 \quad \Rightarrow \quad & \frac{(x - \bar{d}) + (y - \bar{d}) + (z - \bar{d})}{\sigma} + \sqrt{\frac{2}{\pi}} \left\{ \frac{e^{-\frac{1}{2} \left(\frac{x+y-\bar{d}}{\sigma} \right)^2} - e^{-\frac{1}{2} \left(\frac{|x-y|-\bar{d}}{\sigma} \right)^2}}{\operatorname{erf} \left(\frac{x+y-\bar{d}}{\sigma\sqrt{2}} \right) - \operatorname{erf} \left(\frac{|x-y|-\bar{d}}{\sigma\sqrt{2}} \right)} \right. \\ & \left. + (q-1) \frac{e^{-\frac{1}{2} \left(\frac{y-\bar{d}}{\sigma} \right)^2}}{1 - \operatorname{erf} \left(\frac{y-\bar{d}}{\sigma\sqrt{2}} \right)} \right\} = 0, \end{aligned} \quad (50)$$

$$\begin{aligned} \frac{\partial \bar{s}}{\partial \sigma} = 0 \quad \Rightarrow \quad & \frac{(x - \bar{d})^2 + (y - \bar{d})^2 + (z - \bar{d})^2}{\sigma} + \sqrt{\frac{2}{\pi}} \left\{ \frac{(x + y - \bar{d})e^{-\frac{1}{2} \left(\frac{x+y-\bar{d}}{\sigma} \right)^2} - (|x - y| - \bar{d})e^{-\frac{1}{2} \left(\frac{|x-y|-\bar{d}}{\sigma} \right)^2}}{\operatorname{erf} \left(\frac{x+y-\bar{d}}{\sigma\sqrt{2}} \right) - \operatorname{erf} \left(\frac{|x-y|-\bar{d}}{\sigma\sqrt{2}} \right)} \right. \\ & \left. + (q-1) \frac{(y - \bar{d})e^{-\frac{1}{2} \left(\frac{y-\bar{d}}{\sigma} \right)^2}}{1 - \operatorname{erf} \left(\frac{y-\bar{d}}{\sigma\sqrt{2}} \right)} \right\} = 0. \end{aligned} \quad (51)$$

These can be significantly simplified by approximating variables x, y , and z by their means (the ‘‘mean field’’ approximation): $x \sim \bar{x}$, $y \sim \bar{y}$, and $z \sim \bar{z}$. Note that $\bar{x} = \bar{z} = \bar{d}$ and $\bar{y} \leq \bar{d} \Rightarrow |x - y| \sim \bar{d} - \bar{y}$. Introducing a new variable

$$\xi = \frac{\bar{d} - \bar{y}}{\sigma\sqrt{2}}, \quad (52)$$

where \bar{y} is, in fact, a function of \bar{d} and σ , $\bar{y} \equiv \bar{y}(\bar{d}, \sigma)$, we can reduce (50) and (51) to, respectively,

$$\xi - \frac{q-1}{\sqrt{\pi}} \cdot \frac{e^{-\xi^2}}{1 + \operatorname{erf}(\xi)} = 0, \quad (53)$$

$$2\sigma\xi \left\{ \xi - \frac{q-1}{\sqrt{\pi}} \cdot \frac{e^{-\xi^2}}{1 + \operatorname{erf}(\xi)} \right\} + \sqrt{\frac{2}{\pi}} \cdot \frac{\bar{y}e^{-\frac{1}{2} \left(\frac{\bar{y}}{\sigma} \right)^2}}{\operatorname{erf} \left(\frac{\bar{y}}{\sigma\sqrt{2}} \right)} = 0. \quad (54)$$

We can now search for solutions $\xi_{\bar{d}}^*$ of (53) numerically. For $n = 10^4$, $q = \lfloor (n / \log_2 n)^{1/2} \rfloor = 27$, which gives a unique solution $\xi_{\bar{d}}^* = 1.32$.

The direct solution of (54) would involve resolving function $\bar{y}(\bar{d}, \sigma)$ first. However, a simpler way is to use an asymptotic form of the error function in the last term of (54),

$$\operatorname{erf}(\alpha) \sim \frac{2}{\sqrt{\pi}} \alpha e^{-\alpha^2}, \quad \alpha \ll 1, \quad (55)$$

which is valid for sufficiently large σ , $\sigma \gg \bar{y}$. This reduces (54) to

$$2\xi \left[\xi - \frac{q-1}{\sqrt{\pi}} \cdot \frac{e^{-\xi^2}}{1 + \operatorname{erf}(\xi)} \right] + 1 = 0, \quad (56)$$

which has a unique solution $\xi_{\bar{d}}^* = 1.24$.

To see that $\sigma_{\bar{d},\sigma}^*(\bar{d}^*)$ are linear, we just need to check that $\bar{y}(\bar{d},\sigma)$ is a linear function of its arguments in the MSR. Indeed,

$$\bar{y}(\bar{d},\sigma) = \int yg_1(y) dy = \frac{q}{\sigma\sqrt{2\pi}2^{q-1}} \int ye^{-\frac{1}{2}\left(\frac{y-\bar{d}}{\sigma}\right)^2} \left[1 - \operatorname{erf}\left(\frac{y-\bar{d}}{\sigma\sqrt{2}}\right)\right]^{q-1} dy. \quad (57)$$

Changing variables, $\zeta = (y - \bar{d})/(\sigma\sqrt{2})$, and using the asymptotic form of the error function, (55), we see that

$$\bar{y}(\bar{d},\sigma) \sim c_1\sigma + c_2\bar{d}, \quad \text{where} \quad \begin{aligned} c_1 &= \sqrt{\frac{2}{\pi}} \cdot \frac{q}{2^{q-1}} \int \zeta e^{-\zeta^2} \left(1 - \frac{2}{\sqrt{\pi}}\zeta e^{-\zeta^2}\right)^{q-1} d\zeta \\ c_2 &= \frac{1}{\sqrt{\pi}} \cdot \frac{q}{2^{q-1}} \int e^{-\zeta^2} \left(1 - \frac{2}{\sqrt{\pi}}\zeta e^{-\zeta^2}\right)^{q-1} d\zeta \end{aligned}. \quad (58)$$

Substituting this into (52), we find that

$$\sigma_{\bar{d},\sigma}^*(\bar{d}^*) \sim c_{\bar{d},\sigma} \bar{d}^* = \frac{1 - c_2}{\xi_{\bar{d},\sigma}^* \sqrt{2} + c_1} \bar{d}^*. \quad (59)$$

Numerical evaluations of $c_{1,2}$ yield $c_{\bar{d}} = 0.53$ and $c_{\sigma} = 0.57$. We have a good match between $c_{\bar{d}}$ and the data fit (cf. Fig. 3(d)), 0.57. The match between c_{σ} and the data fit, 0.79, is worse, which suggests that the source of error is mostly in (55). We do not obtain non-zero values for the additive coefficients, $\tilde{c}_{\bar{d},\sigma}$, in $\sigma_{\bar{d},\sigma}^*(\bar{d}^*) \sim c_{\bar{d},\sigma} \bar{d}^* + \tilde{c}_{\bar{d},\sigma}$ that define the apex location. Thus, we conclude that a more accurate analysis of the essentially discrete case with small σ is required to analytically obtain the apex location.

Note, however, that equations (53) and (54) are consistent with the observed analytical structure of $\bar{s}(\bar{d},\sigma)$ even for small σ . Indeed, the solutions of the *system* of equations (53) and (54) (corresponding to the true stationary points of \bar{s} , $d\bar{s} = 0$) exist only for $\sigma \rightarrow 0$, when the last term of (54) goes away. Then we have from (52) with $\xi = \xi_{\bar{d}}^*$ that $\bar{y} \rightarrow \bar{d}$ as expected, and *any* \bar{d} delivers a solution. As discussed in Section 3, the actual average distance can be only an integer k or $k + 1/2$ when $\sigma \rightarrow 0$. Thus, the flat-topped peaks and narrow cracks we observe in Fig. 3(c) at $\sigma = 0$ are consistent with $d\bar{s} = 0$ there.

Since $d\bar{s} = 0$ only when $\sigma \rightarrow 0$, we have also effectively demonstrated, at least with the approximations we have made to obtain (53) and (54), that the apex *cannot* be a *true* stationary point of \bar{s} .

The LS size q is a function of n , and, hence, solutions $\xi_{\bar{d},\sigma}^*$ are also functions of the graph size. They are shown in Fig. 6. We see that $\xi_{\bar{d},\sigma}^*(n) = \Theta(\log n)$. This sheds some light on the analytical structure of \bar{s} as a function of n . We can see, for example, that as the network grows, the MSR becomes narrower and closer to $\sigma = 0$. Since for scale-free nets, σ does approach 0 as $n \rightarrow \infty$ ([26]), we conclude that, independent of their size, the scale-free graphs are characterized by the lowest possible average stretch values.

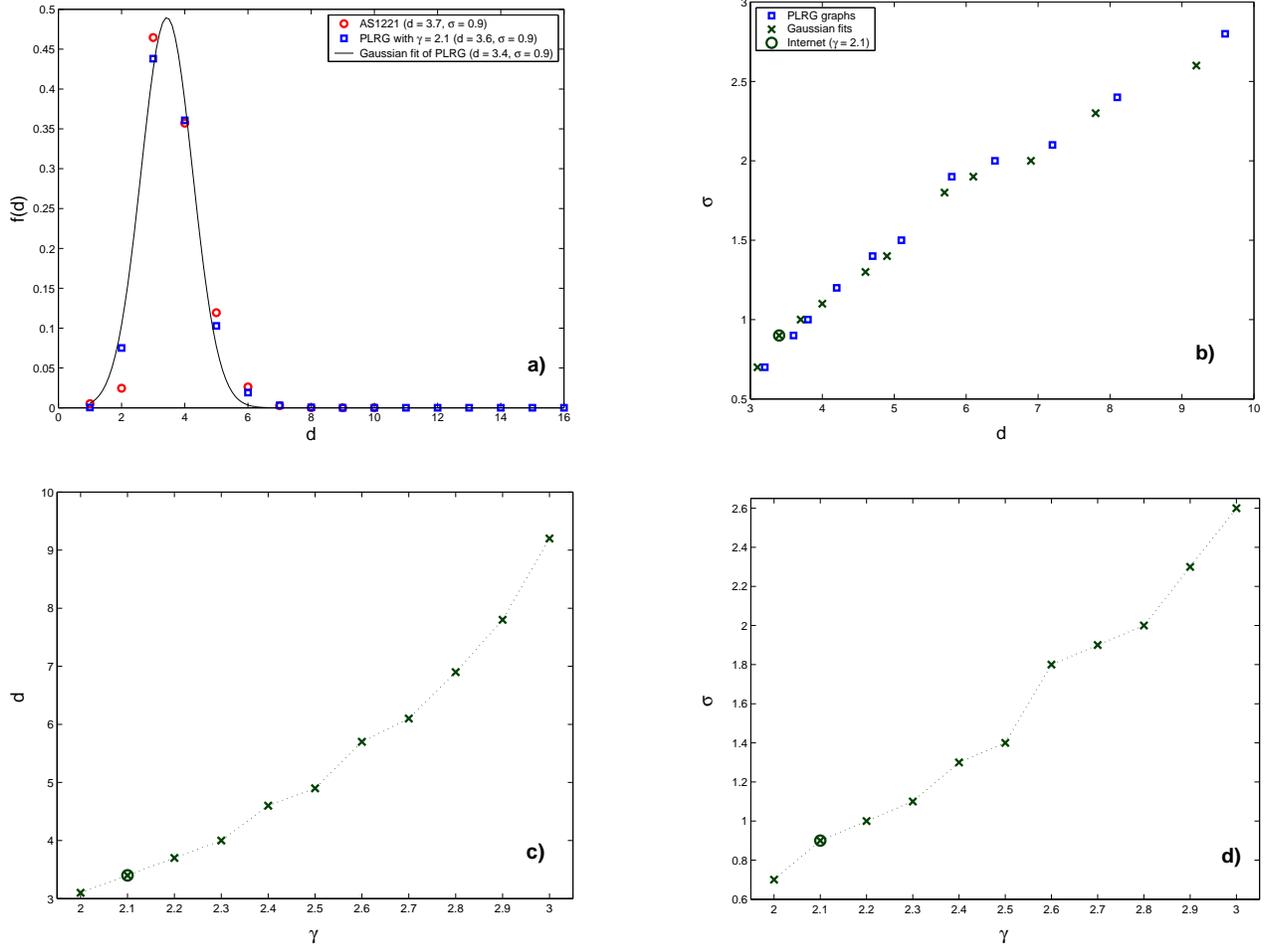


Figure 1: **a)** The distance distributions. The red circles represent the distance distribution from a typical AS (AS #1221) averaged over a period of approximately March-May, 2003. (The source of data is [9]; for other measurements, see [89, 14].) The mean and the standard deviation is 3.7 and 0.9 respectively. The distance distribution in PLRG-generated graphs with $\gamma = 2.1$ is shown by blue squares. The standard deviation is the same as before, the mean is 3.6. The solid line is the Gaussian fit of the PLRG distribution, $\bar{d} = 3.4$ and $\sigma = 0.9$. **b)** The means and standard deviations (blue squares) of distance distributions in PLRG-generated graphs with $\gamma = 2.0, 2.1, \dots, 3.0$ (from left to right), and the corresponding values of \bar{d} and σ (green crosses) in their Gaussian fits. The fitted values of \bar{d} and σ as functions of γ are shown in **(c)** and **(d)** respectively. The Internet value of $\gamma = 2.1$ is circled in (b)-(d).

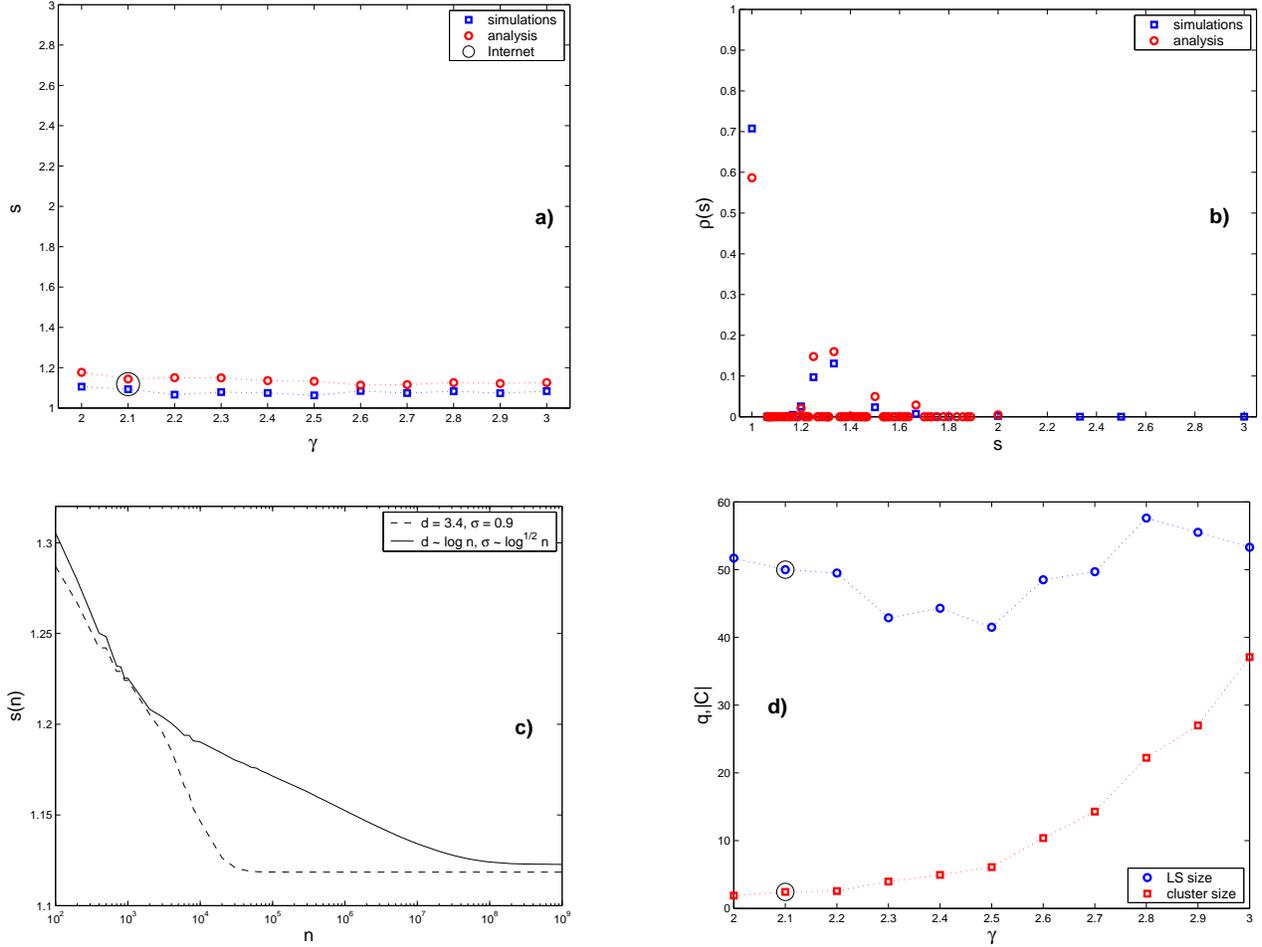


Figure 2: **a)** The analytical results (red circles) and simulation data (blue squares) for the average TZ stretch as a function of γ . **b)** The analytical results (red circles) and simulation data (blue squares) for the TZ stretch distribution with $\gamma = 2.1$. **c)** The analytical data for the average stretch as a function of the graph size. The dashed line corresponds to the case when the distance distribution parameters \bar{d} and σ are fixed to the values observed in the Internet. The solid line presents the data when \bar{d} and σ scale according to the DGM model. **d)** The simulation data for the LS (blue circles) and cluster (red squares) sizes. In the Internet case, $\gamma = 2.1$, the average graph size in simulations is 10,687, the average LS size is 50.0, and the average cluster size is 2.43.

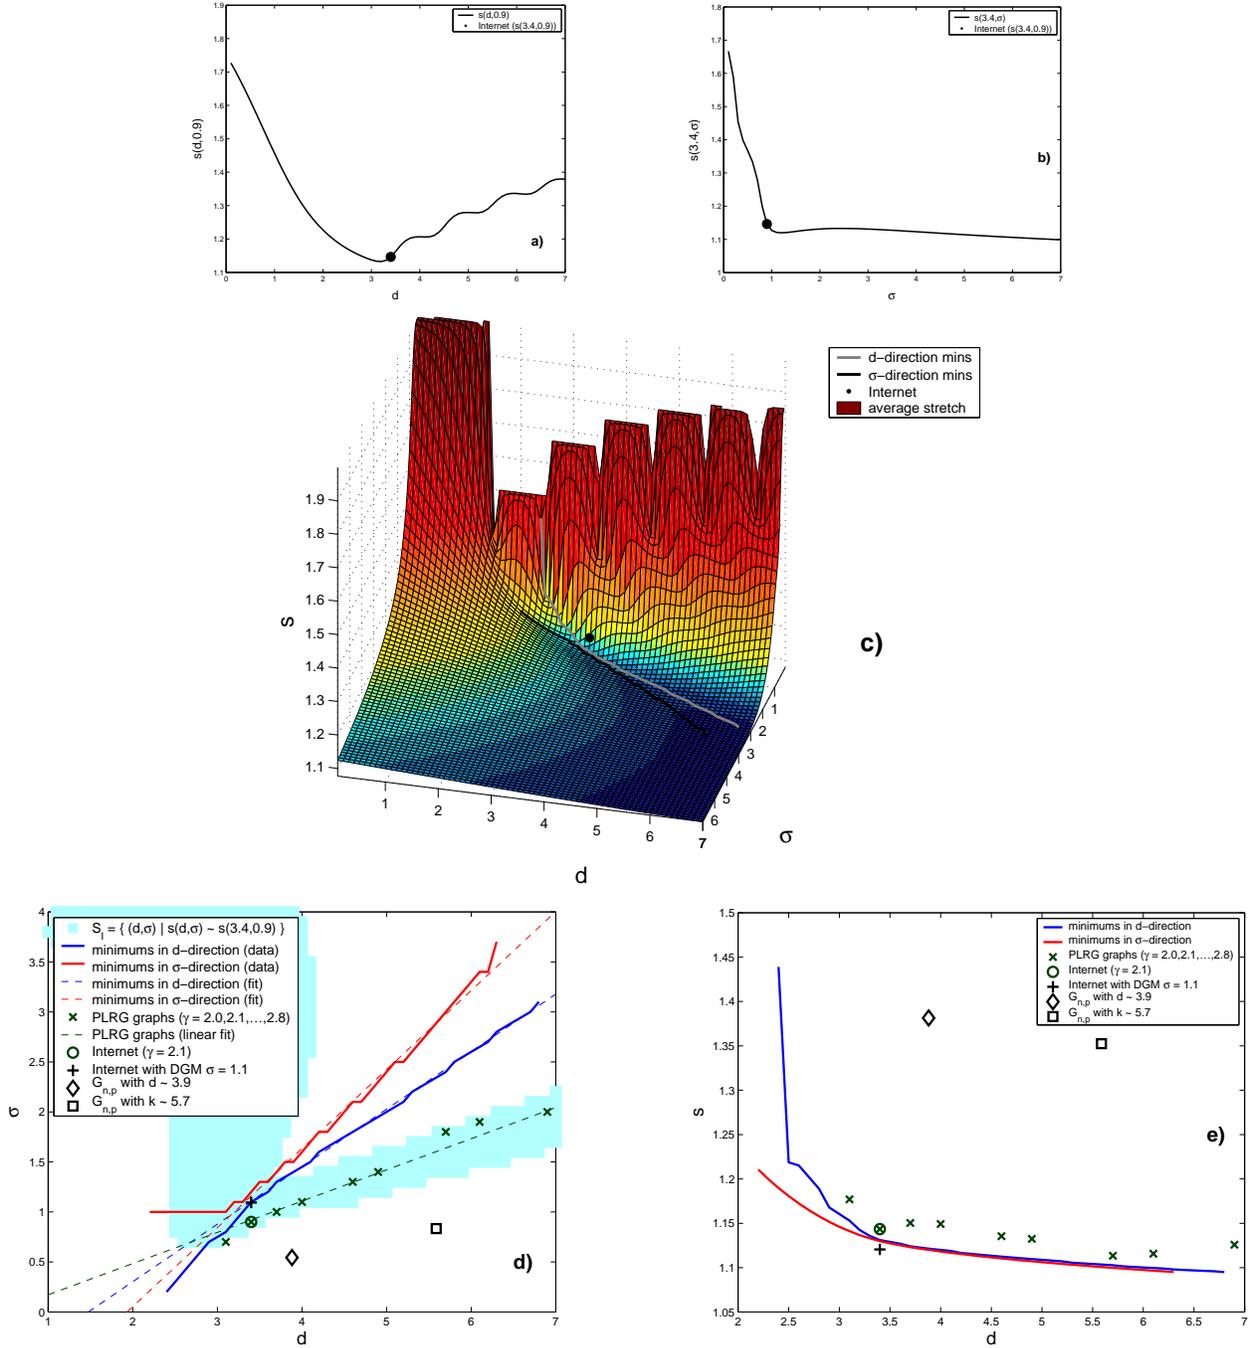


Figure 3: **a),b)** The average stretch as functions of \bar{d} with $\sigma = 0.9$ and of σ with $\bar{d} = 3.4$ respectively. The Internet is represented by the black dot. **c)** The average stretch as a function of \bar{d} and σ . The Internet is represented by the black dot. The stretch minimums along the \bar{d} - and σ -axes, $\sigma_d^*(\bar{d}^*)$ and $\sigma_\sigma^*(\bar{d}^*)$, are the grey and black lines respectively. **d)** The projection of (c) onto the \bar{d} - σ plane. The solid blue (bottom) and red (top) lines represent respectively $\sigma_d^*(\bar{d}^*)$ and $\sigma_\sigma^*(\bar{d}^*)$ (the grey and black lines from (c)). The dashed blue and red lines are their linear fits in the *MSR*. The green crosses are the same as in Fig. 1(b), the green dashed line being their linear fit. The Internet, $\gamma = 2.1$, is circled. The shaded cyan area is S_I from the text. The black plus is the point with the average distance observed in the Internet and the Gaussian width predicted by the DGM model, $\bar{d} = 3.4$, $\sigma = 1.1$. The black diamond and square are the distance distributions of the $G_{n,p}$ graphs from Table 3 matching the Internet average distance and node degree. **e)** The projection of (c) onto the \bar{d} - s plane. The notations are the same as in (d). The graph sizes $n \sim 10^4$ everywhere.

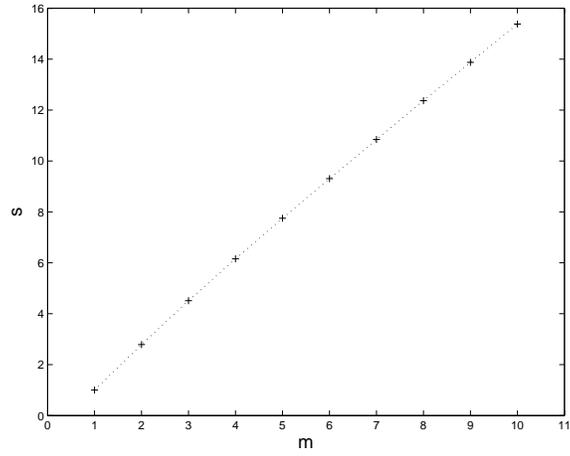


Figure 4: The KK stretch factor s as a function of the number of levels of hierarchy m .

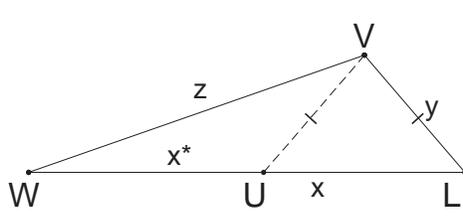


Figure 5: The shortcut triangles.

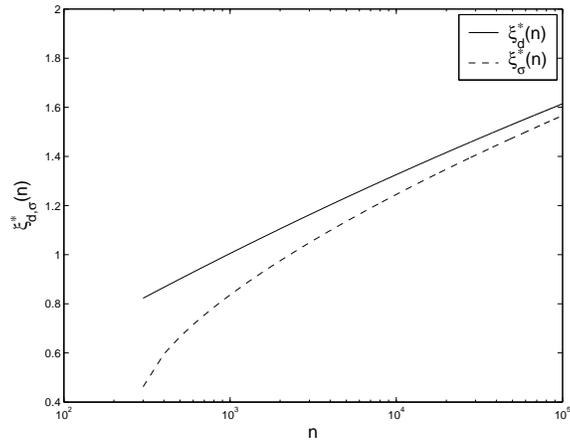


Figure 6: Solutions of equations (53) and (56), ξ_d^* (solid line) and ξ_σ^* (dashed line), as functions of the graph size n .