Conductance and Congestion in Power Law Graphs

Christos Gkantsidis
christos@cc.gatech.edu

Milena Mihail
mihail@cc.gatech.edu

Amin Saberi
saberi@cc.gatech.edu

College of Computing
Georgia Institute of Technology
Atlanta, GA, USA

ABSTRACT

It has been observed that the degrees of the topologies of several communication networks follow heavy tailed statistics. What is the impact of such heavy tailed statistics on the performance of basic communication tasks that a network is presumed to support? How does performance scale with the size of the network? We study routing in families of sparse random graphs whose degrees follow heavy tailed distributions. Instantiations of such random graphs have been proposed as models for the topology of the Internet at the level of Autonomous Systems as well as at the level of routers. Let n be the number of nodes. Suppose that for each pair of nodes with degrees d_u and d_v we have $O(d_u d_v)$ units of demand. Thus the total demand is $O(n^2)$. We argue analytically and experimentally that in the considered random graph model such demand patterns can be routed so that the flow through each link is at most $O\left(n \log^2 n\right)$. This is to be compared with a bound $\Omega\left(n^2\right)$ that holds for arbitrary graphs. Similar results were previously known for sparse random regular graphs, a.k.a. "expander graphs." The significance is that Internet-like topologies, which grow in a dynamic, decentralized fashion and appear highly inhomogeneous, can support routing with performance characteristics comparable to those of their regular counterparts, at least under the assumption of uniform demand and capacities. Our proof uses approximation algorithms for multicommodity flow and establishes strong bounds of a generalization of "expansion," namely "conductance." Besides routing, our bounds on conductance have further implications, most notably on the gap between first and second eigenvalues of the stochastic normalization of the adjacency matrix of the graph.

SIGMETRICS'03, June 10–14, 2003, San Diego, California, USA.

Categories and Subject Descriptors

G.2.2 [**Graph Theory**]: Network Problems; C.2.1 [**Network Architecture and Design**]: Network topology; G.2.1 [**Combinatorics**]: Counting Problems

General Terms

Algorithms, Performance, Theory.

Keywords

Internet topology, power-law graphs, routing, congestion, expansion, conductance.

1. INTRODUCTION

By the mid 90's, when the exponential growth of the Internet became apparent, several predicted an imminent Internet collapse, among other reasons, due to lack of capacity. Recall the memorable quote: "The current onslaught of new users will be more than the Internet infrastructure can support. More new capacity will be needed than the companies that build that infrastructure can deploy" [38, 50]. Not only did this not happen [8], but experimental work suggested that the Internet is performing well with respect to many statistics [50]. The question is whether it was good engineering, inherent properties of the Internet, or, just luck that allowed it to survive this rapid growth without serious performance deterioration. In this paper we argue that the Internet's topology has inherent structural properties that support routing with near-optimal congestion. In particular, we argue that a general family of "Power-Law Random Graphs" (PLRG) [9, 3], of which the Internet topology is believed to be an instantiation [21, 27, 3, 54], has good "expansion" properties. Consequently, approximation algorithms for multicommodity flow imply that this family of graphs supports routing with near-optimal congestion [57, Chapter 21].

Performance is a term that means different things to different parties and different applications. The users of a communication network are concerned with round-trip delay, packet drop probability and throughput. Service providers are concerned with congestion and efficient use of network resources. For applications, such as the WWW, P2P and streaming video, performance metrics become even more context specific. However, there are simple graph-theoretic abstractions believed to capture properties that are essential in "every good network." For example, small diameter (a well defined graph theoretic property) is a desirable characteristic [10, 15], and more recently, the small world phenomenon (a less well defined property, but generally understood to mean small average distance and good clustering) has been identified as a desirable characteristic in searching the WWW and P2P networks [28, 59].

Another graph theoretic abstraction that has been pervasive in algorithms and complexity over the last thirty years is that of *expansion* and its generalizations, namely *cut sparsity* and *conductance* [17, 32, 43, 46, 52, 57]. In a sequence of celebrated results, which generalize the max-flow min-cut theory (e.g. see [34, 35, 57]), cut sparsity and conductance have been explicitly correlated with the performance of routing algorithms as follows. Let $G(V, E, W)$ be an undirected capacitated graph, and let c_e, $e \in E$, denote the capacities. Let $\{(s_1, t_1), \ldots, (s_k, t_k)\}$ be specified pairs of nodes (where each pair is distinct, but a node may be present in several pairs). A separate commodity i is defined for each (s_i, t_i) pair, and for each commodity i a nonnegative demand $\mathrm{dem}(i)$ is also specified. For some routing of this demand, let $l_e \cdot c_e$ denote the flow through link e. The *congestion* of link e is the quantity l_e. Thus the *maximum link congestion* according to this routing is $L = \max_{e\in E} l_e$. The objective is to find a routing that minimizes the maximum link congestion. For example, if we were given the network's topology and demand, L would be the amount of provisioning in the thickest link.

When the routing respects capacity constraints, then l_e is the usual notion of utilization. We shall also consider routings that violate capacity constraints. What is then the meaning of l_e and L? If all demands were scaled by a factor $1/L$, then the capacities would be respected. (The quantity $1/l_e$ is also referred to as "throughput.") Thus, in analyzing L we are making the assumption that $1/L$ is indicative of the worst performance over all links.

Consider a cut $(S, \bar{S})$. Let $\delta(S)$ denote the links that have one endpoint in S and another endpoint in $\bar{S}$. Let $c(S)$ denote the total capacity of the links in the cut: $c(S) = \sum_{e\in\delta(S)} c_e$. Let $\mathrm{dem}(S)$ denote the total demand separated by the cut: $\mathrm{dem}(S) = \sum_{i:|\{s_i,t_i\}\cap S|=1} \mathrm{dem}(i)$. A natural lower bound on L follows by simple averaging principles. In particular, for every routing,

$$\sum_{e\in\delta(S)} l_e c_e \;\geq\; \mathrm{dem}(S)\;,\;\forall\, S \subset V\;.$$

Thus,

$$L \;\geq\; \max_{S\subset V} \frac{\sum_{e\in\delta(S)} l_e c_e}{\sum_{e\in\delta(S)} c_e} \geq \max_{S\subset V} \frac{\mathrm{dem}(S)}{c(S)} = \frac{1}{\min_{S\subset V} \frac{c(S)}{\mathrm{dem}(S)}}.$$

The theory of maximum multicommodity flow developed over the last decade [57, Chapter 21] suggests that there exists a routing with maximum link congestion within a factor $O(\log k)$ of the above lower bound. Moreover, there are known polynomial time algorithms for finding such a routing:

$$\frac{O(\log k)}{\min_{S\subset V} \frac{c(S)}{\mathrm{dem}(S)}} \geq L \geq \frac{1}{\min_{S\subset V} \frac{c(S)}{\mathrm{dem}(S)}}\;. \tag{1}$$

The *cut sparsity* associated with a specific demand is the crucial ratio

$$\min_{S\subset V} \frac{c(S)}{\mathrm{dem}(S)}\;. \tag{2}$$

Let $(d_1, d_2, \ldots, d_n)$ be the degrees of the graph. Let $D = \sum_{i=1}^n d_i = O(n)$ representing the fact that graph is sparse. Consider demand $O(d_u d_v)$ between all pairs of nodes u and v. This includes one unit of demand between all n^2 pairs of nodes as a special case. Define the volume of a set of vertices $S \subset V$ as $\mathrm{vol}(S) = \sum_{u\in S} d_u$. Then (1) implies (details in Theorem 3.1

$$\frac{O(n \log n)}{\Phi} \;\geq\; L \;\geq\; \frac{O(n)}{\Phi} \tag{3}$$

where Φ is the *conductance* and is defined as:

$$\Phi = \min_{S\subset V, \mathrm{vol}(S)\leq D/2} \frac{c(S)}{\sum_{u\in S} d_u} = \min_{S\subset V, \mathrm{vol}(S)\leq D/2} \frac{\sum_{e\in\delta(S)} c_e}{\sum_{u\in S} d_u}. \tag{4}$$

Expander graphs are families of regular graphs (all nodes have the same degree) with linear number of unit capacity links ($|E|$ can be as low as $3|V|$) for which the so-called *expansion factor* is constant:

$$\min_{S\subset V, |S|\leq|V|/2} \frac{\delta(S)}{|S|} \;=\; \Omega(1)\;.$$

For such graphs, and for one unit of demand between all n^2 pairs of nodes, (3) and (4) imply that there exists a polynomial time routing algorithm with maximum link congestion $O(n \log n)$. Alternatively, if all links have capacity $O(n \log n)$ then the demand can be routed with maximum link congestion bounded away from 1. This is optimal since, by simple counting considerations, most of the n^2 pairs have hop distance $\Omega(\log n)$. In a rather strong sense, expander graphs enable excellent resource allocation, since with a linear number of links they support routing with optimal congestion. Note that in arbitrary graphs routing one unit of traffic between every pair of nodes may result in congestion as bad as $\Omega(n^2)$. For example, in a complete binary tree each link incident to the root needs to carry flow $(n/2)^2$. In a tree grown with preferential connectivity there are links incident to the highest degree node that need to carry flow $\Omega(n^{\frac{3}{2}})$.

Admittedly, the known polynomial time algorithms that achieve provably optimal performance (via LP-duality and metric embeddings) are complex and involve non-integral flows [57, Chapter 21]. However, there are complementary results suggesting that near-optimal congestion (up to poly $\log n$ factors) can also be achieved with integral short paths and decentralized, on-line algorithms (e.g. see [23, 29].) Therefore, a constant expansion factor is thought of as an "excellent promise" for routing. Random regular graphs are long known to possess constant expansion [43, Chapter 5.6]. These, together with explicit constructions [36] have found many applications in networks: [47, 48, 49] for non-blocking networks, [32] for parallel architectures, [33] for circuit switching, [31] for peer-to-peer networks, to list just a handful that span three decades. All these applications involve expanders as carefully constructed mathematical and engineering artifacts.

The networking paradigm is shifting. Today's open, distributed and dynamic networks are no longer artifacts that we construct, but phenomena that we study. One of the first observed striking differences is in the distribution of the degrees of the underlying network topologies: The degrees of the WWW [6, 12], the Internet at the level of Autonomous Systems [21] and at the router level [54], and several other examples [54], all follow heavy-tailed statistics, often expressed as *power-laws*: The frequency of nodes with degree d is proportional to $d^{-\zeta}$, for some constant ζ typically between 2 and 3. At the same time, these remain sparse, linear-size networks (the average degree is around 4 for Internet topologies and around 8 for the WWW.) The main result of this paper is:

> *Power Law Random Graphs (PLRG) can support routing of $O(d_u d_v)$ units of flow between each pair of vertices u and v with degrees d_u and d_v respectively, with congestion $O\left(n \log^2 n\right)$. This includes unit demand between all pairs of nodes as a special case.*

This is only a $\log n$ factor off from the congestion achieved by linear size regular graphs with constant expansion. Thus our result can be understood as follows. The skewed degree distributions of PLRGs result in a hierarchical organization, with nodes of (typically) high degree forming the "core" of the network [53, 54, 26]. This is reminiscent of a tree-like structure. Intuitively, we expect that links in the core carry more flow. Our result suggests that the bound $O(n \log^2 n)$ by which the flow scales in the core of PLRGs is closer to the bound $O(n \log n)$ of a robust flat structure, such as an expander, rather than the bound $\Omega(n^{1+\epsilon})$ of a tree.

What is the implication of our result for real networks, like the Internet, that are believed to be instantiations of PLRGs? Even though our model is very simple to capture the full complexity of real systems, we believe that it carries a positive message. We view the moderate rate at which the congestion grows and the established strong conductance as an indication that, despite its decentralized uncoordinated dynamic growth, the Internet preserves good resource allocation and load balancing properties and prevents extreme fragilities and monopolies [56, 19, 58]. Admittedly, this is a subjective statement, but we believe that it is a firm starting point.

In summary, we have analyzed routing on PLRG's under the assumptions:

(a) All links have the same capacity.

(b) Demand $O(d_u d_v)$ between all pairs of nodes (unit uniform demand is a special case.)

(c) The objective is to minimize maximum link congestion.

(d) Flows can be fractional and involve paths of arbitrary length.

We made assumptions (a) and (b) due to lack of publicly available information about link capacities and demand patterns. For practical purposes and for the worst case analysis considered here, we believe that assumption (a) is not particularly restrictive. See the definition of conductance (4). Intuitively, we expect that unbalance in link capacities will favor links belonging to cuts for which $\frac{|\delta(S)|}{\sum_{u \in S} d_u}$ is small. That only helps Φ.

We believe that assumption (b) is quite restrictive. In particular, it does not capture popular sites that may have low degree, but high demand. However, the notion of cut sparsity can be defined for arbitrary capacities and demands and the general methodology of (1) and (2) carries over. Thus the methodology of our work provides a starting point to study more general demand patterns, once such patterns become better characterized.

We believe that assumption (c) is not particularly restrictive, though this is also subjective. Assumption (c) essentially imposes worst case analysis. There are many other performance metrics in computer science theory and networking, but worst case analysis is a reasonable place to start.

Assumptions (c) and (d) imply that the objective is to minimize the maximum link congestion. In many networks of practical interest the objectives can be different. For example many protocols minimize the hop count for every source-destination pair. We have experimentally measured the congestion in power-law random graphs under shortest-hop integral routing (in a set-up reminiscent of the Internet at the level of Autonomous Systems.) Our measurements indicate that the congestion still increases like $O(n \text{poly} \log n)$.

A key technical ingredient in our proofs is to show that the core of the PLRG has strong conductance properties. These have further implications, most notably on the spectral gap of the stochastic normalization of the adjacency matrix of the core of the graph. Such spectral gaps have found in the past many algorithmic applications [17, 52, 43, 57]. For example, they imply reliability and fast cover times and hitting times —the latter are related to crawling and searching. Thus, we believe that our bounds on conductance and spectral gap will be a useful tool in establishing further properties for Internet performance. (Similar bounds have been subsequently obtained for the model of growth with preferential attachment [40]). In passing, we also note that the rather sharp bounds obtained in Corollary 3.4 are further evidence that the spectrum of the stochastic normalization of the adjacency matrix is an important metric for Internet topologies (see also [21, 39, 26]).

The balance of the paper is as follows:

In Section 2 we discuss random graph models for graphs with skewed degrees, including Internet topologies, and formalize the random model that is suitable for our study.

In Section 3 we give a theoretical argument based on conductance and along the lines of (3) and (4) that shows that maximum link congestion is $O(n \log^2 n)$. This section contains the conductance and eigenvalue separation proofs. The analytical argument applies to random graphs under certain model restrictions. We view these restrictions as mild, but certainly, in a strict sense, they do not include the whole class of power-law random graphs. More importantly, by invoking approximation algorithms for multicommodity flows, they involve non-integral flows and centralized routing along paths that are not necessarily short.

In Section 4 we validate the $O(n \text{poly} \log n)$ congestion bound experimentally, for integral routing along shortest paths using as graphs real and synthetic Internet topologies. We further compare the congestion of Internet and Internet-like topologies to trees and 3-regular expanders. These are worst-case and best-case sparse random graphs with respect to congestion. The congestion in real and synthetic Internet topologies appears to scale in a way much closer to that of expanders than trees.

Summary and open problems are in Section 5.

2. STRUCTURAL MODELS FOR GRAPHS WITH SKEWED DEGREE SEQUENCES

Random graph models producing graphs with skewed degree sequences fall into two general categories: evolutionary and structural. *Evolutionary* models identify growth primitives giving rise to skewed degree distributions. These primitives can be microscopic, such as multiobjective optimization [20, 5], and macroscopic, such as statistical preferential connectivity [6, 18, 37, 30]. The advantage of microscopic evolutionary models is that they may capture additional network semantics. Their disadvantage is that they are hard to simulate and notoriously hard to analyze (e.g. see [20].) This is due to the detailed optimization problems solved in each step, and the dependencies between steps; dependencies pose the biggest hurdle in probabilistic arguments. The advantage of macroscopic evolutionary models is that they are easy to simulate. Their disadvantage is that they are also quite difficult to analyze, due, again, to the dependencies between steps (e.g. see [11].) *Structural* models start with a given skewed degree distribution, perhaps a power-law predicting the degrees of a real network [1, 27], and interpolate a graph that matches the degree sequence (exactly or approximately) and satisfies certain other randomness properties [27, 3, 55, 25]. A big advantage of such structural models is their amenability to analytical treatment. By taking the degree sequence as a granted, most of the dependencies arising in the analysis of evolutionary models can be removed [3, 16, 39]. This has been also noted by mathematicians who have known several robustness properties in structural random graph models for some time [9, 41, 42], though the term used there is *configurational*. In addition, structural models have been found good fits for Internet topologies [54]. Therefore, in this paper where we attempt analytical justifications, we will use a structural model similar to the one in [3].

Let $\vec{d} = (d_1, d_2, \dots, d_n)$ be a sequence of integers. The structural or configurational method generates a random graph as follows. First consider $D = \sum_{i=1}^{n} d_i$ mini-vertices; think of mini-vertices as lying in n *clusters* of size d_i, $1 \le i \le n$. Then construct a random perfect matching among the mini-vertices and generate a graph on the n original vertices as suggested by this perfect matching in the natural way: two vertices are connected with an edge if and only if at least one edge of the random perfect matching was connecting mini-vertices of their corresponding clusters [9, 4]. This is an uncapacited graph. Alternatively, we may generate a capacitated graph by assigning capacity c_e between vertices u and v proportional to the number of edges between the clusters of d_u and d_v mini-vertices corresponding to u and v. Note that this is a general random graph model. It makes no assumptions on the starting sequence of integers.

A power-law random graph (PLRG) is an uncapacitated graph generated according to the structural method for a degree sequence obtained by sampling from a power-law distribution. In [4] it was shown mathematically that a PLRG consists of more than one connected components, almost surely, though it has a giant connected component, almost surely. Necessary and sufficient conditions under which a general degree sequence results in a connected graph in the structural model were obtained in [41, 42]. From the technical perspective, notice that it might be hard to argue about expansion on a random graph model that is not even connected. The intuition is that, what causes small isolated connected components in the entire graph, may cause small sets with bad expansion inside the giant component.

In [54] it was argued experimentally that the giant component of a PLRG matches several characteristics of real complex networks, and hence is a good candidate for generating synthetic Internet topologies. However, one notable discrepancy between PLRGs and topologies of real communications networks is that in real networks nodes of very small degree (customers) are much more likely to be connected to nodes of large degree (providers). This has been measured in [14, 44] and has been formalized in [44].

We will use a technical modification of PLRG that ensures connectivity, almost surely, and always connects nodes of degree 1 and 2 to nodes of degree greater than 3. In practice, we will construct a modified PLRG as follows. For a degree sequence $\vec{d} = (d_1, d_2, \dots, d_n)$, we first consider a connected graph that satisfies the degree sequence exactly and is reminiscent of some Internet topology. For example, we may consider the graph generated by Inet [27], any connected graph that satisfies the degree sequence and some further randomness criterion [54, 7], or a Markov chain Monte Carlo simulation approach [25]. We perform iterated pruning of vertices of degrees 1 and 2, until we are left with a graph whose smallest degree is 3. (The significance of "3" is that this is the smallest constant for which random graphs are connected expanders; for example, random 3-regular graphs are almost surely connected expanders, while random 2-regular graphs are almost surely disjoint cycles.) Let $\vec{\delta} = (\delta_1, \delta_2, \dots, \delta_n)$ be the degree sequence of the graph after the pruning. We consider a PLRG generated for the non-zero degree vertices of $\vec{\delta}$; we call this PLRG a *core*. Note that cores can be parameterized according to the smallest degree vertex that they contain; this natural notion has been repeatedly observed whenever iterated pruning has been considered (e.g. see [13].) Finally, we attach the pruned vertices with their original degrees in a preferential fashion. This can be done as follows. For each vertex u in the core consider $d_u - \delta_u$ mini-vertices. Let U be the set of mini-vertices arising from vertices in the core. For each pruned vertex v consider d_v mini-vertices. Let V be the set of mini-vertices arising from pruned vertices. We may now construct a random maximum matching between U and V and connect u to v if and only if some mini-vertex arising from u was connected to some mini-vertex arising from v.

Ideally, the modified PLRG model described above can become completely formal, once the degrees $\vec{\delta}$ of pruned Internet topologies are characterized. We have measured these degrees in a method similar to [21] and have found them to also obey well characterizable heavy tailed statistics (this is not surprising, neither intuitively nor analytically.) In this write-up we refrain from further descriptions of these tedious but straightforward measurements, in particular, because they are not necessary for the analytic argument. Indeed, the analytic argument holds for any sequence of integers $\vec{d} = (d_1, d_2, \dots, d_n)$ with

$$D = \sum_{i=1}^{n} d_i = O(n) \text{ and } \max\{d_i,\ 1 \le i \le n\} = O(n^{\frac{1}{2}}) \quad (5)$$

and any sequence $\vec{\delta} = (\delta_1, \delta_2, \dots, \delta_n)$ with

$$\begin{array}{rcll} \delta_i & \le & d_i & \forall i \\ \sum_{i:\delta_i=0} d_i & \le & \sum_{i:\delta_i\neq 0} d_i - \delta_i & \\ \{\delta_i \ : \ \delta_i > 0\}| & = & \Omega(n) & \\ \delta_i & = & \Omega(d_i) & \forall i \ : \ \delta_i > 0 \end{array} \quad (6)$$

Conditions (6) are straightforward. They say that there are enough degrees in the core to absorb the non-core vertices, that the core is a constant fraction of the entire network and that the degree of each vertex inside the core is proportional to its degree in the entire network.

Conditions (5) say that the network is sparse and that the maximum degree is $O(n^{\frac{1}{2}})$. We justify the latter by pointing out that it is true in the evolutionary model of growth with preferential connectivity and the structural PLRG. For the evolutionary model it was proved in [10]. For the structural model it can be verified as follows. Consider any degree sequence produced by n independent samples drawn from a power law distribution with $\Pr[d_i = d] \simeq d^{-\zeta}$, $2 < \zeta < 3$. The probability that any drawn sample has value greater than $\Omega(n^{\epsilon})$ is $O(n^{-\epsilon\zeta})$. Thus the probability that all drawn samples are smaller than $O(n^{\epsilon})$ is $O(n^{1-\epsilon\zeta}) = o(1)$, for $\epsilon > \frac{1}{\zeta}$. Since $2 < \zeta < 3$, $\epsilon = \frac{1}{2}$ suffices.

3. THEORETICAL SUPPORT: THE CONDUCTANCE ARGUMENT

Theorem 3.1. *Let $\vec{d}$ and $\vec{\delta}$ be sequences of integers satisfying (5) and (6). Suppose that all links have unit capacity. Then, there exists a way to route demand $O(d_u d_v)$ between all pairs of nodes u and v such that all links have flow $O(n \log^2 n)$. Moreover, the routing can be computed in polynomial time.*

Proof. Every node that does not belong to the core, that is where $\delta_u = 0$, will transfer its demand uniformly to the core vertices to which it is attached. Because of (6), this imposes demand $O(\delta_u \delta_v)$ between all pairs of vertices in the core.

To argue about the routing of the demand in the core we may start from (2). In particular, where V' is the set of vertices belonging to the core ($V' = \{u \in V : \delta_u > 0\}$), the cut sparsity in the core is:

$$\begin{aligned} \min_{S\subset V'} \frac{c(S)}{\mathrm{dem}(S)} &\geq \Omega\left(\min_{S\subset V'} \frac{c(S)}{\sum_{u\in S}\delta_u \sum_{v\in\bar{S}}\delta_v}\right) \\ &\geq \Omega\left(\min_{S\subset V', \mathrm{vol}(S)\leq \mathrm{vol}(\bar{S})} \frac{c(S)}{\Omega(n)\sum_{u\in S} d_u}\right) \\ &= \frac{\Phi}{\Omega(n)} . \end{aligned}$$

From Lemma 3.3, for the core, we have $\Phi = \Omega(1/\log n)$. Thus the cut sparsity in the core is $\Omega(1/n \log n)$. Now (1) implies that there is a routing of all the demands with maximum link flow $O(n \log^2 n)$. $\square$

We proceed to establish conductance for the core. This is done in Lemma 3.3. The main technical step in in Lemma 3.2 that follows.

Lemma 3.2. *[Main Lemma.] Let $\vec{d} = d_1 \geq d_2 \geq \ldots \geq d_n$ be a sequence of integers with*

$$d_n \geq d_{\min} = 3 \quad \text{and} \quad D = \sum_{i=1}^{n} d_i = O(n) .$$

Let $G(V, E, W)$ be a graph with capacities generated according to the structural random graph model of Section 2. The conductance of $G(V, E, W)$ is

$$\Phi(G) = \min_{S\subset V, \mathrm{vol}(S)\leq D/2} \frac{c(S)}{\mathrm{vol}(S)} \;\geq\; \Omega(1) ,$$

with probability $1 - o(1)$.

Proof. For a positive constant α, we say that a set of vertices S with $k = \mathrm{vol}(S) \leq D/2$ is Bad if

$$\frac{c(S)}{\mathrm{vol}(S)} < \alpha n .$$

We will show that there exists a positive constant α such that

$$\Pr[\exists\ \mathsf{Bad}S] \;\leq\; o(1) . \tag{7}$$

The left hand side of (7) is

$$\sum_{k=d_{\min}}^{D/2} \Pr[\exists\ \mathsf{Bad}S,\ \mathrm{vol}(S) = k] . \tag{8}$$

Let us fix k in the above range. There are at most $\binom{D/d_{\min}}{k/d_{\min}}$ sets of vertices in G that have volume k. This is because every such set arises from a set of mini-vertices such that the total number of minivertices is k and, for each cluster, either all or none of the mini-vertices of the cluster are included. Since the minimum cluster size is $d_{\min}$, the number of possibilities is maximized if all clusters were of size $d_{\min}$. We hence need to bound

$$\sum_{k=d_{\min}}^{D/2} \binom{D/d_{\min}}{k/d_{\min}} \Pr[\text{a fixed set } S, \mathrm{vol}(S) = k, \text{ is } \mathsf{Bad}] . \tag{9}$$

We analyze the probabilities that appear in the terms of the above summation. We may now assume that the set S is fixed. Let A denote the set of the k mini-vertices corresponding to S. Let $\bar{A}$ denote the set of the $(D - k)$ mini-vertices corresponding to $\bar{S}$. Let $B_A \subset A$ be the set of mini-vertices in A that were matched to mini-vertices in $\bar{A}$. Let $B_{\bar{A}} \subset \bar{A}$ be the set of mini-vertices in $\bar{A}$ that were matched to mini-vertices in A. In order for S to be Bad, the cardinality $|B_A| = |B_{\bar{A}}|$ is at most αk. For each cardinality in the range 0 to αk, there are at most $\binom{k}{\alpha k}$ ways to fix the mini-vertices in B_A and at most $\binom{D-k}{\alpha k}$ ways to fix the mini-vertices in $B_{\bar{A}}$. We may now assume that the sets B_A and $B_{\bar{A}}$ are fixed. We need to analyze the probability that the random perfect matching on the mini-vertices matched all mini-vertices in $A \setminus B_A$ inside $A \setminus B_A$, all mini-vertices in $\bar{A} \setminus B_{\bar{A}}$ inside $\bar{A} \setminus B_{\bar{A}}$, and all vertices in $B_A \cup B_{\bar{A}}$ inside $B_A \cup B_{\bar{A}}$. The above probability can be expressed in terms of the total number of perfect matchings on n vertices. Let $f(n) = \frac{n!}{2^{n/2}\cdot(n/2)!}$ be the number of perfect matchings on n vertices. We may now write

$$\begin{gathered} \Pr[\text{a fixed set } S, \mathrm{vol}(S) = k, \text{ is } \mathsf{Bad}] \\ \leq \\ \alpha k \binom{k}{\alpha k}\binom{D-k}{\alpha k} \frac{f(2\alpha k)\, f(k-\alpha k)\, f(D-k-\alpha k)}{f(D)} \end{gathered} \tag{10}$$

We proceed with calculations. We bound each term of (10) separately. We will repeatedly use the following bounds which follow from Stirling's approximation [9, p.4]. There are positive constants c_1 and c_2 such that,

$$c_1\, n^{n+1/2}\, e^{-n} \;<\; n! \;<\; c_2\, n^{n+1/2}\, e^{-n}. \tag{11}$$

Using $\binom{n}{m} = \frac{n!}{m!\,(n-m)!}$ and (11), for some constant c_3, we can bound

$$\binom{D/d_{\min}}{k/d_{\min}} < c_3 \left(\frac{D^D}{k^k\ (D-k)^{D-k}}\right)^{1/d_{\min}} . \tag{12}$$

Using the inequality $\binom{n}{m} \leq (n\,e/m)^m$, we bound

$$\binom{k}{\alpha k}\binom{D-k}{\alpha k} \leq \left(\frac{e}{\alpha}\right)^{2\alpha k}\left(\frac{D-k}{k}\right)^{\alpha k} \tag{13}$$

By substituting the values of f and using (11), for some constant c_4, we can bound

$$\begin{aligned}
\alpha k\,\frac{f(\alpha k)\,f(k-\alpha k)\,f(D-k-\alpha k)}{f(D)} &= \\
\alpha k\,\frac{(2\alpha k)!}{2^{\alpha k}(\alpha k)!}\,\frac{(k-\alpha k)!}{2^{k(1-\alpha)/2}(k(1-\alpha)/2)!}\cdot & \\
\frac{(D-(1+\alpha)k)!}{2^{\frac{D-(1+\alpha)k}{2}}\left(\frac{D-(1+\alpha)k}{2}\right)!}\,\frac{2^{D/2}(D/2)!}{D!} &< \\
c_3\,\alpha k\frac{2^{2\alpha k}2^{1/2}(\alpha k)^{2\alpha k}}{(\alpha k)^{\alpha k}}\,\frac{(k(1-\alpha))^{k(1-\alpha)}\,2^{k(1-\alpha)/2}2^{1/2}}{(k(1-\alpha))^{k(1-\alpha)/2}}\cdot & \\
\frac{(D-(1+\alpha)k)^{D-(1+\alpha)k}\,2^{\frac{D-(1+\alpha)k}{2}}}{(D-(1+\alpha)k)^{\frac{D-(1+\alpha)k}{2}}}\,\frac{D^{D/2}}{D^D\,2^{D/2}} &< \\
c_4\,\alpha k\,(2\alpha)^{\alpha k}\,\frac{k^{k(1+\alpha)/2}\,(D-k)^{(D-k-\alpha k)/2}}{D^{D/2}} &.
\end{aligned} \tag{14}$$

Now combining (12), (13) and (14) we get that, for some constant c_5,

$$\begin{gathered}
\binom{D/d_{\min}}{k/d_{\min}}\binom{k}{\alpha k}\binom{D-k}{\alpha k}\alpha k\,\frac{f(\alpha k)\,f(k-\alpha k)\,f(D-k-\alpha k)}{f(D)} \\
< \\
c_5\alpha k\frac{2^{\alpha k}\,e^{2\alpha k}}{\alpha^{\alpha k}}\,\left(\frac{k}{D}\right)^{k((1-\alpha)/2-1/d_{\min})}
\end{gathered} \tag{15}$$

Define

$$\beta = \frac{2e^2}{\alpha} \quad\text{and}\quad \gamma = \frac{1-\alpha}{2} - \frac{1}{d_{\min}}\ . \tag{16}$$

Define

$$G\left(k\right) = c_5\alpha k\,\beta^{\alpha k}\left(\frac{k}{D}\right)^{\gamma k}\ . \tag{17}$$

Using (10), (15), (16) and (17), we can bound the quantity in (9) by

$$\sum_{k=d_{\min}}^{D/2} G\left(k\right)$$

Note that a necessary condition for the sum to be bounded is that the terms $G(k)$ become vanishing, which, by (17), requires $\gamma > 0$, which, by (16) implies $d_{\min} > \frac{2}{1-\alpha} \Rightarrow d_{\min} \geq 3$ and $\alpha \leq \alpha_1(d_{\min}) = 1 - \frac{2}{d_{\min}}$.

The first derivative of $G(k)$ is

$$\frac{d\ G(k)}{d\ k} = \left(\frac{1}{k} + \alpha\ln\beta + \gamma\ln\frac{k}{D} + \gamma\right)G(k) \tag{18}$$

The second derivative of $G(k)$ is

$$\frac{d^2\ G(k)}{d\ k^2} = \left(-\frac{1}{k^2} + \frac{\gamma}{k} + \left(\frac{d\ G(k)}{d\ k}\right)^2\right)G(k) \tag{19}$$

The first derivative is negative for $k = 3d_{\min}$ and sufficiently large D. The second derivative is positive for $k \geq 3d_{\min}$ and $\alpha \leq \alpha_2(d_{\min}) = 1 - \frac{8}{3d_{\min}}$. Notice that $d_{\min} \geq 3$ guarantees that α is positive. Thus $G(k)$ attains its maximum either at $G(3d_{\min})$ or at $G(D/2)$. We wish to bound both these quantities by $o(1/D)$.

$$G(3d_{\min}) = c_6\frac{1}{D^{3d_{\min}\gamma}} \leq o\left(\frac{1}{D}\right)$$

for constant c_6 and $3d_{\min}\gamma > 1$ which implies $\alpha \leq \alpha_2(d_{\min})$.

$$G(D/2) = c_7\frac{D}{2}\left(\frac{\beta^\alpha}{2^\gamma}\right)^{D/2} \leq o\left(\frac{1}{D}\right)$$

for constant c_7 and

$$\beta^\alpha < 2^\gamma \Rightarrow \alpha\left(\frac{3}{2} + 2\log_2 e - \log_2\alpha\right) < \frac{1}{2} - \frac{1}{d_{\min}}\ .$$

The left side of the inequality is a monotonically increasing function of α in the range $(0,1]$. A sufficient condition for the inequality to hold gives the third condition for α: $\alpha < \alpha_3(d_{\min})$ ($\alpha < 0.0175$ for $d_{\min} = 3$ suffices).

For k in the range $d_{\min}$ to $3d_{\min}$, there are a constant numbers of terms of (17). It can be seen that each one of these terms is $o(1)$.

Thus:

$$\begin{aligned}
\sum_{k=d_{\min}}^{D/2} G(k) &= \sum_{k=d_{\min}}^{3d_{\min}-1} G(k) + \sum_{k=3d_{\min}}^{D/2} G(k) \\
&\leq o(1) + D\cdot o(1/D) \\
&= o(1)
\end{aligned}$$

This completes the proof of (7) by using (9), (15) and the definition of $G(k)$ (17). □

LEMMA 3.3. *Let $\vec{d} = d_1 \geq d_2 \geq \ldots \geq d_n$ be a sequence of integers with*

$$d_1 = O(n^{\frac{1}{2}}),\quad d_n \geq 3 \quad\text{and}\quad D = \sum_{i=1}^{n} d_i = O(n)\ . \tag{20}$$

Let $G(V,E,W)$ be a graph with capacities generated according to the random graph model of Section 2. Let $G(V,E)$ be the corresponding uncapacitated random graph. The conductance of $G(V,E)$ is

$$\Phi(G) \geq \Omega(1/\log n)\ ,$$

with probability $1 - o(1)$.

PROOF. In Lemma 3.2 we showed conductance $\Omega(1)$ for the capacitated $G(V,E,W)$. Therefore it suffices to show that no link will have capacity more than $O(\log n)$, almost surely. In turn, it suffices to bound the probability that the link between vertices u and v of degrees d_1 and d_2 respectively have capacity more than $O(\log n)$. We will bound this probability by $o(1/n^2)$. This probability is maximized when $d_1 = d_2 = \Theta(n^{\frac{1}{2}})$.

Let u_1, …, u_{d_1} be the mini-vertices corresponding to u. Let Y_{u_i}, $1 \leq i \leq d_1$, be 1 if u_i is connected to a mini-vertex corresponding to the cluster of minivertices of v and 0 otherwise. We are interested in the capacity $Y = \sum_{i=1}^{d_1} Y_{u_i}$. This can be bounded by the sum X of $d_1 = \Theta(\sqrt{n})$ independent Bernoulli trials, each one with probability of success $\frac{d_1}{D-d_1} = \Theta(1/\sqrt{n})$. The expectation $\mu = E[X] = \Theta(1)$.

Using the standard tail equality (Chernoff bound) is [9, p.12]:

$$\Pr[X > \mu + t] \;<\; e^{-\frac{t^2}{2(\mu+t/3)}}$$

we get

$$\Pr[Y > \Omega(\log n)] = o(1/n^2) \;.$$

□

COROLLARY 3.4. *Let $G(V, E)$ be a random graph as in Lemma 3.3. Let A be the adjacency matrix of G. Consider a stochastic matrix P corresponding to a random walk G. The spectrum of P is in $[-1, 1]$, with the largest eigenvalue $\lambda_1 = 1$. Let λ_2 be the second largest eigenvalue. Then,*

$$1 - \Omega\left(\frac{1}{\log n}\right) \;<\; \lambda_2 \;<\; 1 - \Omega\left(\frac{1}{\log^2 n}\right) \;.$$

PROOF. Follows from the known inequalities (e.g. see [52, p.53])

$$1 - 2\Phi \;<\; \lambda_2 \;<\; 1 - \frac{\Phi^2}{2}$$

and Lemma 3.2. □

4. EVALUATION

4.1 Methodology

In the previous sections we proved congestion properties for routing involving non-integral flows over paths of arbitrary length. In this section, we strive to experimentally verify that the basic conclusions hold even under shortest-hop routing involving integral flows.

A canonical example of particular interest is Internet routing at the level of Autonomous Systems (AS). The routing protocol at the AS level is the Border Gateway Protocol (BGP) [51]. This protocol filters all paths that do not satisfy certain policy constraints and among the remaining ones, it picks paths with minimum hop distance. The policy constraints can be arbitrarily complex, but at minimum they are set in such a way to prevent transit traffic [24]. Network owners are willing to route traffic that originates or terminates in their own network or in the networks of their (paying) customers. They do not allow traffic arriving from one of their (non-paying) peers to be forwarded to another peer, since this wastes network resources without generating revenues.

We model two routing schemes. The first one is *Shortest-Hop Routing.* This scheme routes along paths with minimum hop (AS) distance. In case of multiple minimum hop paths, we pick one of them at random. The second is *Shortest-Hop Routing with Policies,* which is an extension of the previous scheme without transit traffic. The first scheme can be applied to all topologies. The second scheme can be used only when information about the type of the links is available, which is the case for the real topologies and not for the synthetic ones.

As in the theoretical analysis, we have assumed that there is one unit of demand between any two ASes, which is routed over a single shortest paths. For each link we compute the number of paths going through that link (link congestion), and we examine the maximum number over all links (congestion). We also examine the profile of these values. We study how the maximum link congestion evolves as the size of the network increases and how the rate of increase compares to two baseline models. The first baseline model is the family of 3-regular expanders. Graphs of this type have congestion $O(n \log n)$. The second baseline model is the family of trees grown with preferential connectivity. These graphs have congestion $\Omega(n^{3/2})$. To see this, we use the fact that such trees have a node of degree $\Omega(n^{1/2})$ [10]. Let u be this node. Let v be a neighbor of u such that the subtree rooted at v has $\Omega(n^{1/2})$. Let K be this subtree. If $|K| < \frac{n}{2}$, then the complement of K has at least $n/2$ nodes. Thus, the link (u, v) carries flow $n^{1/2}\frac{n}{2}$. If $|K| \geq \frac{n}{2}$, then the link (u, v) must carry flow between K and the neighbors of u outside K. This is $\frac{n}{2}(n^{1/2} - 1)$. We find qualitatively that the congestion of the AS topology is closer to 3-regular expanders than trees grown with preferential connectivity.

In Section 4.2 we discuss the data used. In Section 4.3 we discuss the evolution of congestion with time. Further observations are in Section 4.4.

4.2 Data Used

We have used AS topology data from two sources. The first source is [2]. In addition to the topology, [2] classifies links as customer-provider, or peer-to-peer. This classification is important for policy routing. We have data from [2] for the years 2001 and 2002.

The second set of data is [22]. Though this set is far less complete, it has the advantage that it spans the time period of 1997 to 2001. [22] does not contain information about the relationships between the ASes. We have used the algorithm of [24] to infer AS relationships.

Using these data to derive conclusive results raises two issues. First, we have no way of knowing how complete the measurements are. Second, in order to study the evolution of the Internet from 1997 to 2002, we had to rely on data from two different sources which have different levels of accuracy [14]. Thus, the data are not directly comparable and we can observe minor irregularities in the results. Nevertheless, we believe that the trends we observe are correct.

To study the evolution of the Internet in the future we need to experiment with topologies that are larger than the current Internet. We use the modified structural random graph model described in Section 2 to generate Internet-like topologies for large sizes with degree sequences obtained from Inet [27].

Note that there is an alternative representation of the Internet at the level of routers. The resulting graph contains much more detailed connectivity information. Even though the AS graph is a less detailed representation of the Internet than the router-level graph, it has three main advantages. First, it is smaller and thus amenable to processing. Second, detailed data for the AS topology are collected since 1997 [22]. No router level data that span that many years exist. Third, bottlenecks in the Internet are usually either at the access links, or at the connections between ASes (on the contrary, links inside the ASes are usually overprovisioned) [45].

The construction of trees with preferential connectivity and 3-regular expanders is straightforward. Our 3-regular expanders of size n are random graphs with n nodes, where

Table 1: Congestion of the AS topology

Year	Nodes	Links	Shortest Hop	Shortest Hop with Policies
1997	3055	5238	559653	1213810
1998	4341	7949	681067	2184841
1999	6209	12206	1051782	3914276
2000	9318	18272	1754897	8376841
2001	10915	23757	5523085	7533275
2002	13155	28317	9096654	11361893

Note: Maximum congestion over all links for different instances of the AS topology, assuming min-hop routing and min-hop routing with policies.

each node has degree 3 (technically this is achieved by superpositioning three random perfect matchings.)

4.3 Evolution of Congestion with Time

First, we examine the evolution of the maximum link congestion over all links for the AS topology. The results for various snapshots of the Internet topology are given in Table 1. The number of nodes increased by a factor of 4.3 between 1997 and 2002. The congestion increased by a factor of 13.36 and 9.36 for graphs with shortest hop routing without and with policies respectively. These results agree qualitatively with Section 3, which argues that congestion increases as a function of $n \cdot \text{poly} \log n$, where n is the number of nodes.

The transition from year 2000 to 2001 in Table 1 may come as an anomaly since the congestion for shortest hop routing with policies decreased. We believe that this is due to the fact that we used data collected from different sources.

To study the evolution of the congestion as the size of the network increases, we need to experiment with graphs that have a wider range in sizes. We have used the Inet generator to get an initial degree sequence $\vec{d}$. We used the Markov-Chain method to get a random graph for $\vec{d}$ [25]. We used the method of Section 2 to get a degree sequence $\vec{\delta}$ for the core. We used the modified PLRG model of Section 2 to get the final topology Figure 1 gives the maximum link congestion for graphs of size 3037, 5000, 7000, …, 23000. The observation for this figure is the shape of the curve. The values of the congestion are bounded from above and below by $O\left(n \cdot \log^{2.6} n\right)$ and $O\left(n \cdot \log^{2} n\right)$ respectively. These bounds were not designed to be tight, but to illustrate that congestion in Internet-like topologies grows as a function of $O\left(n \cdot \log^{k} n\right)$ for a small value of k. The observations remain the same if we use the output of Inet, or the unmodified PLRG.

How does the congestion of synthetic PLRG compares to that of 3-regular expanders and trees? We give this comparison in Figure 2. The modified PLRG model appears to behave much closer to the 3-regular expander than the tree. The evolution of congestion for 3-regular expanders and trees are given in Figures 3 and 4 respectively.

4.4 Congestion Fingerprints

Next, we examine in more detail the characteristics of the link congestion for all links. Consider all the links of the graph sorted according to their congestion in increasing order. We call the resulting profile the *congestion fingerprint* of the graph.

The congestion fingerprint of the AS topology is drawn in

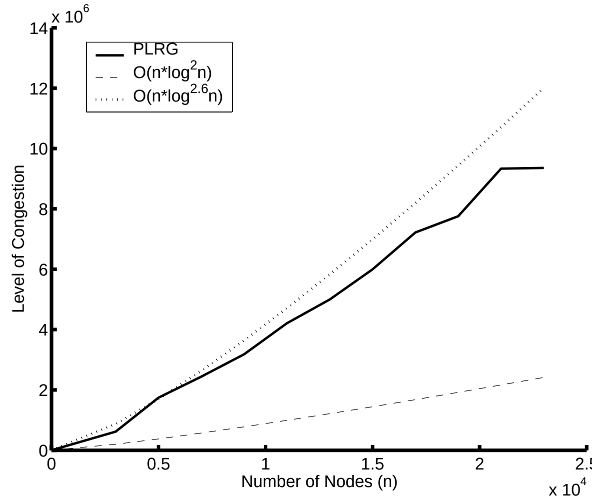

Figure 1: Congestion for PLRG topologies.

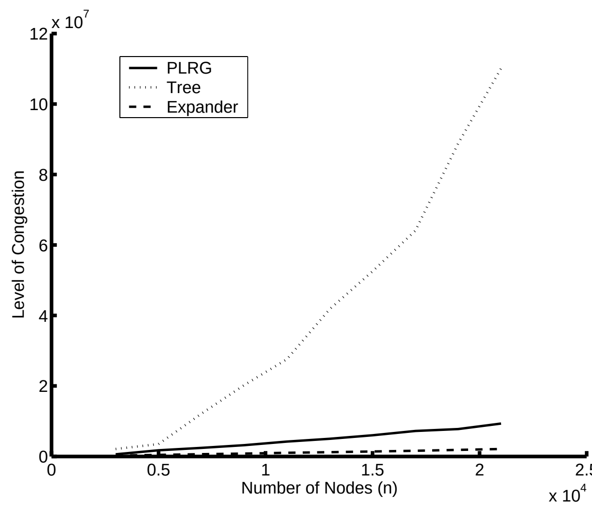

Figure 2: Comparison of congestion between synthetic PLRGs, 3-regular expanders and trees grown with preferential connectivity.

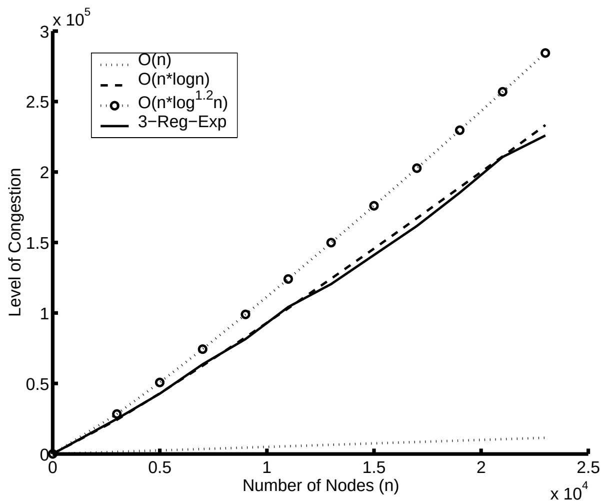

Figure 3: Congestion for 3-regular expanders.

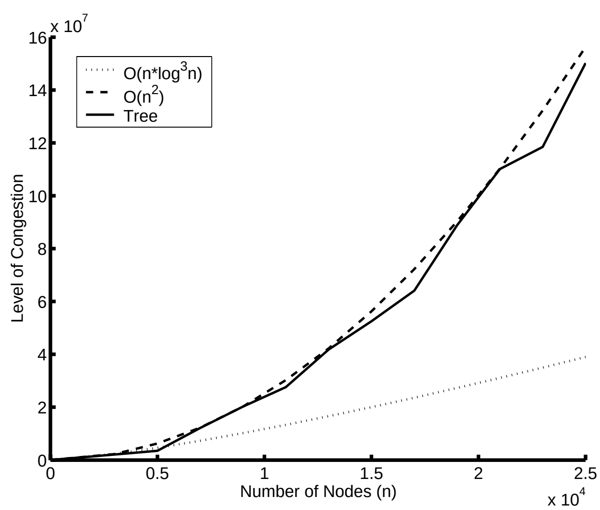

Figure 4: Congestion for tree grown with preferential connectivity.

Table 2: Degrees of endpoints for the ten most congested links.

Congestion	Degree 1	Degree 2
9096654	2640	625
4223943	2640	1530
3704681	2640	586
3267646	2640	795
2742594	2640	159
2481002	2640	140
2188507	2640	837
1714812	586	1530
1682390	2640	330
1681173	2640	191

Results for a topology of 13155 nodes using integral shortest hop routing without policies.

Figure 5. The fingerprint is given for various snapshots of the Internet and for the routing with and without policies. Below, we discuss some further observations.

The most congested links have always one endpoint (and very often both) of high degree (see Table 2.) Nodes with high degrees correspond to big providers. Thus, congestion appears around the core.

The difference between the most congested links and the average congestion over all links is 2-3 orders of magnitude. The difference between the most congested links and the least congested ones is 7 orders of magnitude. Observe that the maximum possible difference is bounded above by n^2, which is less than 210^8 for graphs with $n < 14K$.

Around 10% of the links are heavily congested. These are the peaks of Figure 5. In this sense, our paper can be understood as quantifying $O\left(n \cdot \text{poly}\log n\right)$ as the growth rate of the "peak."

Both fingerprints show the same trends. This means that the use of policy routing does not give extra information for our purposes. Thus, we believe that the analysis of non-policy routing in the synthetic topologies approximates well routing with policies for our purposes.

Using the degree sequences of real topologies, we generated synthetic topologies using the PLRG and the modified PLRG. The congestion fingerprints of the generated graphs are qualitatively similar to the real topologies and very different of that of expanders (see Figures 6 and 7.)

5. SUMMARY

In this paper we studied the impact of the Internet topology on the performance of basic communications tasks, such as routing with low congestion. We particularly focused on the issue of scaling. We proposed to use the theory of approximation algorithms for maximum multicommodity flow to bound congestion [57].

For power law random graphs we showed congestion scaling as $O(n \log^2 n)$, where n is the number of nodes, under uniform capacity and demand. This is near-optimal and places power law random graphs much closer to expander graphs than trees. Power law random graphs are believed to be good models for Internet topologies.

Using our methodology to establish results for more general capacity and demand patterns is an interesting open problem.

The above results concern non-integral routing along paths

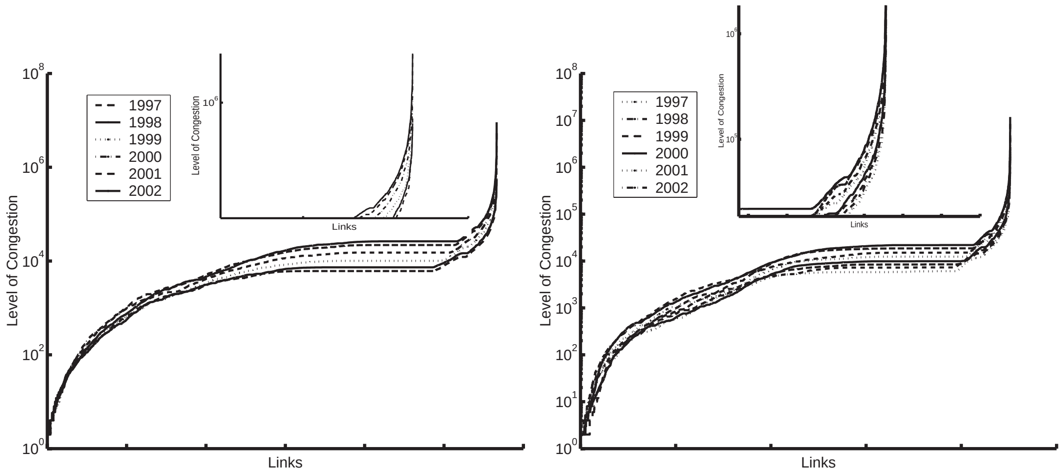

Figure 5: Congestion Fingerprint: Congestion of all links for various snapshots of the AS topology assuming min-hop routing (left) and min-hop routing with policies (right). Links are sorted from the least congested to the most. Observe the log scale in the vertical axis. We have stretched the curves horizontally to be comparable in size. This is important since different topologies have different number of edges.

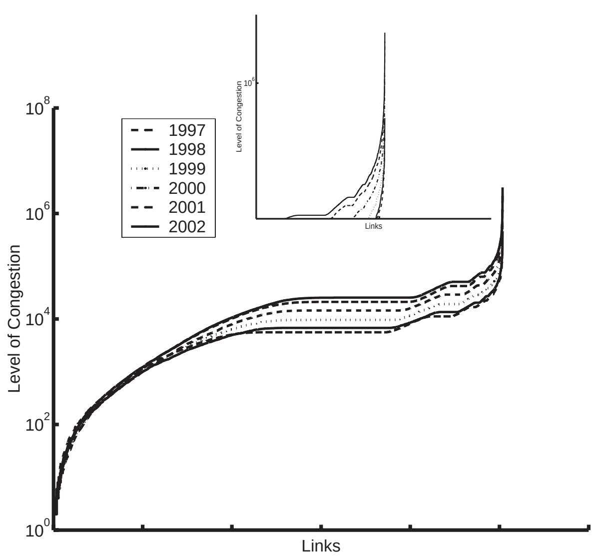

Figure 6: Congestion fingerprint for PLRG.

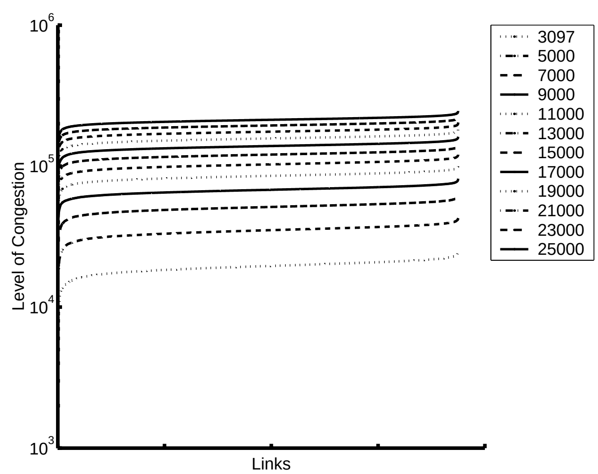

Figure 7: Congestion fingerprints of the 3-regular expander topologies (vertical axis in log scale.)

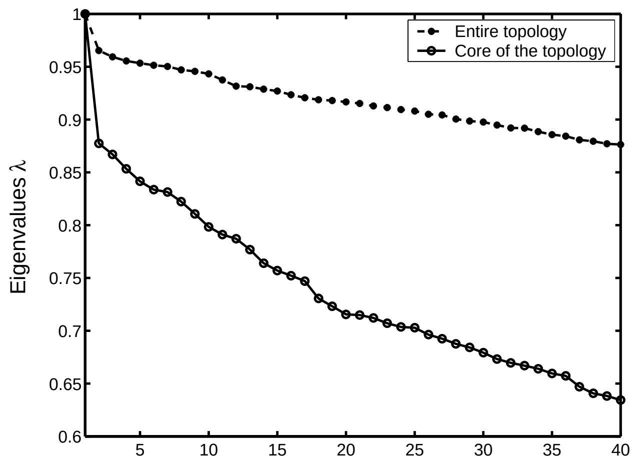

Figure 8: Eigenvalues of the Internet topology: The biggest eigenvalues of the stochastic matrix arising from the Internet at the AS level, for both the core and the entire topology [25]. Note the gap between the first and second eigenvalue in the core of the topology and see Corollary 3.4.

of arbitrary length. We gave experimental evidence that the results hold even for routing along integral shortest paths. We use data from the Internet at the level of Autonomous Systems (AS). This level implements the BGP routing protocol.

The key ingredient of our technical argument establishes strong conductance properties in the core of Power-Law Random Graphs. The bound of conductance is of independent interest. In particular, it implies strong bounds on the second eigenvalue of the stochastic normalization of the adjacency matrix. This also agrees with experimental findings, see Figure 8. Such spectral gaps have wide implications in algorithms and stochastic processes [17, 52, 43]. Exploring further implications for the Internet performance is an interesting open problem.

6. ACKNOWLEDGMENTS

Special thanks to Christos Papadimitriou for bringing to our attention the question of congestion in the core of the Internet. Thanks to Ashish Goel and the anonymous reviewers for their comments on the presentation. The authors were supported by NSF under ITR-0220343.

7. REFERENCES

[1] L. Adamic. Zipf, power-laws, and pareto, a ranking tutorial. *http://www.hpl.hp.com/shl/papers/ranking/*, 2002.

[2] S. Agarwal, L. Subramanian, J. Rexford, and R. Katz. Characterizing the internet hierarchy from multiple vantage points web page. *http://www.cs.berkeley.edu/~sagarwal/research/BGP-hierarchy/*, 2002.

[3] W. Aiello, F. R. K. Chung, and L. Lu. A random graph model for power law graphs. In *Proc. 41st Symposium on Foundations of Computer Science (FOCS)*, pages 171–180. IEEE, 2000.

[4] W. Aiello, F. R. K. Chung, and L. Lu. Random evolution in massive graphs. In *Proc. 42nd Symposium on Foundations of Computer Science (FOCS)*, pages 510–519. IEEE, 2001.

[5] D. Alderson, J. Doyle, and W. Willinger. Toward an optimization-driven framework for designing and generating realistic internet topologies. In *HotNets*, 2002.

[6] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[7] C. Berge. *Graphs and Hypergraphs*. North Holland Publishing Company, 1973.

[8] Bob metcalfe eats his words. *Internet Computing Online*, 1(3), 1997. http://www.computer.org/internet/v1n3/eats9702.htm.

[9] B. Bollobás. *Random Graphs, 2nd Ed.* Cambridge University Press, 2001.

[10] B. Bollobás and O. Riordan. The diameter of scale-free graphs. *Combinatorica*, To appear.

[11] B. Bollobás, O. Riordan, J. Spencer, and G. Tusnády. The degree sequence of a scale-free random graph process. *Random Structures and Algorithms*, 18(3):279–290, 2001.

[12] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomikns, and J. Wiener. Graph structure in the web. *9th International World Wide Web Conference (WWW9)/Computer Networks*, 33(1-6):309–320, 2000.

[13] A. Broido and K. Claffy. Internet topology: Connectivity of ip graphs. http://www.caida.org/outreach/papers/2001/OSD/, 2001.

[14] C. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origins of power-laws in internet topologies revisited. In *Proc. Infocom*. IEEE, 2002.

[15] F. Chung and L. Lu. The average distance in a random graph with given expected degree. Available at http://math.ucsd.edu/~fan/inet.html.

[16] F. Chung and L. Lu. Connected components in random graphs with given degree sequences. *http://www.math.ucsd.edu/~fan*, 2002.

[17] F. R. Chung. *Spectral Graph Theory*. American Mathematical Society Book Series, 1997.

[18] C. Cooper and A. Frieze. A general model for web graphs. In *ESA*, pages 500–511, 2001.

[19] J. Doyle, J. Carlson, S. Low, F. Paganini, G. Vinnicombe, W. Willinger, J. Hickey, P. Parrilo, and L. Vandenberghe. Robustness and the internet: Theoretical foundations. *http://netlab.caltech.edu/internet/*, 2002.

[20] A. Fabrikant, E. Koutsoupias, and C. Papadimitriou. Heuristically optimized tradeoffs: A new paradigm for powerlaws in the internet. *ICALP 2002*, 2002.

[21] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proc. SigComm*. ACM, 1999.

[22] N. L. for Applied Retwork Research. Route views archive. http://moat.nlanr.net/Routing/rawdata, 2002.

[23] A. Frieze. *Disjoint Paths in Expander Graphs via Random Walks: A Short Survey*, volume 1518 of

Lecture Notes in Computer Science. Springer Verlag, 1998.

[24] L. Gao. On inferring autonomous system relationships in the internet. *IEEE Glogal Internet*, 2000.

[25] C. Gkantsidis, M. Mihail, and E. Zegura. The markov chain simulation method for generating connected power law random graphs. In *Proc. 5th Workshop on Algorithm Engineering and Experiments (ALENEX)*. SIAM, January 2003.

[26] C. Gkantsidis, M. Mihail, and E. Zegura. Spectral analysis of internet topologies. In *Proc. Infocom*. IEEE, 2003. To appear.

[27] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR-433-00, U. Michigan, Ann Arbor, 2000.

[28] J. Kleinberg. Navigation in small world. *Nature*, 406, 2000.

[29] J. Kleinberg and R. Rubinfeld. Short paths in expander graphs. In *Proc. 37th Symposium on Foundations of Computer Science (FOCS)*, pages 86–95. IEEE, 1996.

[30] R. Kumar, P. Raghavan, S. Rajagopalan, S. D., A. Tomkins, and E. Upfal. Stochastic models for the web graphs. In *Proc. 41st Symposium on Foundations of Computer Science (FOCS)*, pages 57–65. IEEE, 2000.

[31] C. Law and K.-Y. Siu. Distributed construction of random expander networks. In *Proc. Infocom*. IEEE, 2003. To appear.

[32] F. Leighton. *Parallel Algorithms and Architectures*. Morgan Kaufmann, 1992.

[33] F. Leighton and S. Rao. Circuit switching: Multicommodity flow based approach. In *Workshop on Randomized Parallel Computing*, 1996.

[34] F. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46:787–832, 1999.

[35] N. Linial, E. London, and Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995.

[36] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8:261–278, 1988.

[37] A. Medina, I. Matta, and J. Byers. On the origin of power-laws in internet topologies. *ACM Computer Communications Review*, April 2000.

[38] B. Metcalfe. *Internet Collapses and Other InfoWorld Punditry*. Hungry Minds, Inc., May 2000.

[39] M. Mihail and C. Papadimitriou. *Random 2002*, 2002.

[40] M. Mihail, C. Papadimitrious, and A. Saberi. On certain connectivity properties of internet graphs. *submitted*, 2003.

[41] M. Molloy and B. Reed. A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, 6:161–180, 1995.

[42] M. Molloy and B. Reed. The size of the largest component of a random graph on a fixed degree sequence. *Combinatorics, Probability and Computing*, 7:295–306, 1998.

[43] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[44] M. Newman. Assortativity mixing in networks. *Phys. Rev. Lett.*, 89, 2002.

[45] V. Padmanabhan, L. Qiu, and H. Wang. Server-based inference of internet performance. In *Proc. Infocom*. IEEE, 2003. To appear.

[46] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[47] N. Pippinger. Information theory and the complexity of switching networks. In *Proc. 16th Symposium on Foundations of Computer Science (FOCS)*, pages 113–118, 1975.

[48] N. Pippinger. Superconcentrators. *SIAM Journal on Computing*, 6(2):298–304, 1977.

[49] N. Pippinger. On rearrangeable and non-blocking switching networks. *Journal of Computer and System Sciences*, 17(2):145–162, 1978.

[50] J. S. Quarterman. Imminent death of the internet? *Matrix News*, 6(6), June 1996. Also, available at http://www2.aus.us.mids.org/mn/606/death.html.

[51] Y. Rehkter and P. Gross. Application of the border gateway protocol in the internet. RFC 1655, IETF, July 1994.

[52] A. Sinclair. *Algorithms for Random Generation and Counting: A Markov Chain Approach*. Springer-Verlag, 1997.

[53] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the internet hierarchy from multiple vantage points. In *Proc. Infocom*. IEEE, 2002.

[54] H. Tagmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: Degree-based vs structural. In *Proc. SigComm*. ACM, 2002.

[55] H. Tagmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: Degree-based vs structural. Technical Report 02-760, USC, 2002.

[56] D. Towsley. Modeling the internet: Seeing the forest through the trees. Keynote Address, Sigmetrics, 2002.

[57] V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.

[58] W. Willinger and J. Doyle. Robustness and the internet: Design and evolution. *http://netlab.caltech.edu/internet/*, 2002.

[59] H. Zhang, A. Goel, and R. Govindan. Using the small-world model to improve freenet performance. In *Proc. Infocom*. IEEE, 2002.