On Certain Connectivity Properties of the Internet Topology

Milena Mihail
College of Computing
Georgia Tech.
mihail@cc.gatech.edu

Christos Papadimitriou
Dept. of Computer Science
U.C. Berkeley
christos@cs.berkeley.edu

Amin Saberi
College of Computing
Georgia Tech.
saberi@cc.gatech.edu

June 14, 2004

Abstract

We show that random graphs in the preferential connectivity model have constant conductance, and hence have worst-case routing congestion that scales logarithmically with the number of nodes. Another immediate implication is constant spectral gap between the first and second eigenvalues of the random walk matrix associated with these graphs. We also show that the expected frugality (overpayment in the Vickrey-Clarke-Groves mechanism for shortest paths) of a sparse Erdős-Renyi random graph is bounded by a small constant.

1 Introduction

The Internet is a computational system of immense complexity that was not designed by a single entity, but emerged from the *ad hoc* interactions of many entities on the basis of ground rules that were deliberately open and minimally restrictive. As a result, it is the first computational artifact that must be studied by observation, measurement, and the development and validation of hypotheses, models and falsifiable theories —in a manner not unlike the one in which other sciences approach the universe, the brain, the cell, and the market. This paper aims to contribute to the growing corpus of mathematical results and techniques that are pertinent to this novel, within Computer Science, research mode.

Since connectivity is a network's *raison d' être*, it is no surprise that various aspects of the Internet's connectivity (such as degrees, diameter, cuts, and tolerance to element failures) have been the subject of intense study, measurement, and speculation, see e.g. [6, 24, 46, 14, 10, 4, **?**]. In this paper we address two sophisticated aspects of connectivity that are particularly relevant to the Internet, namely *conductance* and *frugality*.

As the Internet grows, extensive measurements show a clear congestion increase in the core and relate this to network performance (e.g. see [4, 29, 30, 44, 27]). Therefore, one of the most crucial questions one can ask is, *how does the congestion at the Internet's core scale with the number of nodes?* In other words, if we assume unit traffic between all nodes (more accurately, traffic weighted by some measure of the size of each node, typically captured by its degree), how do the loads on the

edges balance? Since the Internet is a very sparse graph (average degree between 3 and 4 [38, 34]), there are two extremes to consider here: In constant degree trees one expects that congestion (traffic in the worst edge) grows as n^2 with the nodes, while in constant-degree expanders this growth is close to the theoretical minimum, $n \log n$.

The observation in [24] that the degree distribution of the Internet has heavy tails, or is "scale-free" (has deviations from the mean that decrease only polynomially, forming a straight line in log-log plot) has brought center-stage several models of random graphs that exhibit such degree distributions; it is thus compelling to estimate the asymptotic growth of congestion in scale-free random graph models. In this paper we consider the model of *growth with preferential attachment* in which an arriving node connects with d edges to previously arrived nodes chosen with probability proportional to the degrees of the latter [6, 32, 3, 11, 19]. We show that, for $d \geq 2$, *almost all scale-free graphs* in this model have *constant conductance*; as a corollary, approximate multicommodity flow algorithms imply routing with congestion $O(n \log n)$. An immediate additional implication is constant spectral gap between the first and second eigenvalues of the stochastic normalization of the adjacency matrix of the graph. This is also in accordance with measurement: [26] found the second eigenvalue of the Internet topology between .8 and .9 (and of its core between .6 and .7) for snapshots between 1997 and 2002 during which the network has grown by a factor of 20. Elsewhere, [7] measure a gap for the (symmetrized, degree-homogenized) graph of the world-wide web, again over a long period of observations.

A persistent technical difficulty in treating graphs grown with preferential connectivity arises from the inhomogeneity and dependencies between edges [32, 10, 11, 23]. The crux of our proof is in establishing a bound that is invariant of time (shifting argument in Lemma 2). Prior to our work, [27] and [18] had shown conductance and spectral gap $\Omega(1/\log n)$ and $\Omega(1/\text{poly}\log n)$ respectively, for structural scale-free random graph models (Erdős-Renyi adaptations for skewed degree sequences). Structural scale-free random graph models avoid all dependencies between vertices, and are hence easier to analyze [2, 17, 36, 18]. However, in those models certain bad events occur almost surely and inverse logarithmic factors appear unavoidable. More relevant to this paper, [20] had shown conductance $\Omega(1/\log n)$ for the growth with preferential connectivity model considered here, and for constant d much larger than 2. In view of the above, our result (Theorem 1) is the first constant characterization of these fundamental measures.

The graph of the Internet autonomous systems (AS's) is formed as these entities enter into service or peering agreements seeking to ensure good connectivity for their customers. The incentive structure of the situation has been the subject of much study, speculation, and mystification. On the other hand, recently we have seen the development of a novel research corpus in the interface between Algorithms and Game Theory, aiming exactly at understanding such incentive issues in connection to the Internet [39, 41, 42]. Already in the pioneering paper by Nisan and Ronen [39], the shortest path problem in a network was studied as an interesting application of mechanism design, and it was shown that the Vickrey-Clarke-Groves (VCG) mechanism [35] indeed yields an incentive-compatible mechanism for routing (roughly speaking, a protocol in which nodes with private cost information are willing to reveal their true costs). This was taken one step further in [25], where it was shown that such mechanisms can in fact be implemented with minimum overhead and disruption on BGP, the currently predominant interdomain routing protocol [28]. On the negative side, it has been observed [5] that the VCG mechanism, as well as any other "reasonable" mechanism, leads to very significant overpayments in the worst case —an unbounded multiplicative factor above the original

cost of the shortest path. (The VCG overpayment of an edge on the shortest path is the amount by which the cost of the shortest path is increased when this edge is deleted, if this amount is finite, and is defined to be zero otherwise. The term "frugality" has been used —as it turns out, a little too optimistically— to denote this quantity.) Very recently, [22] established that this is inherent in any incentive-compatible protocol for finding shortest paths in graphs. Despite these negative worst-case results, however, [25] measure the VCG overpayment in the Internet graph, assuming unit costs, and observe that it is very modest (about .3, or 30%, compared to the unbounded factors predicted in the literature). As the frugality of a network is a compelling metaphor for the "competitiveness" implicit in its topology, the issue is of some importance. *Are there any mathematical reasons to expect that frugality is small on the average in random graphs? Or is the low-overpayment phenomenon observed in [25] evidence of strategic evasion of monopolistic situations by AS's in the Internet?*

In Section 3 of this paper we provide a partial answer, establishing that, with high probability, the expected VCG overpayment, over all origin-destination pairs and all edges in the shortest path, is bounded both from above by a non-increasing function of the expected degree. We show this in the Erdős-Renyi $G_{n,p}$ model when $np = \omega(\log n)$ by analysis of the shortest and second shortest paths in a random graph. Obtaining a similar result for random graphs whose degree sequences follow power-law degree distributions is an interesting open problem.

2 On the Conductance of Scale Free Graphs

In this section we use the notation $G_{d,n}$ to denote graphs grown with preferential attachment. We will use the following definitions for these random graph processes. $G_{1,n} = T_n$ is a tree grown in n time steps, one vertex at each time step. Its vertices are called *mini-vertices* and they are named after the time that they arrive. At time 1 the tree consists of a single mini-vertex with a self loop. At time t, $2 \leq t \leq n$, mini-vertex t arrives and attaches with a single edge to a mini-vertex t', $t' < t$, chosen among all mini-vertices with probability proportional to their degrees at time $t-1$. We call the mini-vertex t' to which mini-vertex t attached the *father* of t (let the father of 1 be 1, by convention). For $d \geq 2$, the graph $G_{d,n}$ is generated by first growing a tree T_{dn} and then, for $1 \leq \tau \leq n$, contracting mini-vertices $d\tau - i$, for $0 \leq i \leq d-1$. Self-loops and multiple edges are preserved. We call the vertex of $G_{d,n}$ that resulted by contracting mini-vertices $d\tau - i$, $0 \leq i \leq d-1$, vertex τ. Thus, for every $S \subset [n]$, we may associate a subset of vertices of the graph $G_{d,n}$ and a subset of mini-vertices of the tree T_{dn} in the natural way: mini-vertex $d\tau - i$, $0 \leq i \leq d-1$, is associated with S if and only if $\tau \in S$.

Let $G(V, E)$ be an undirected multigraph with self-loops. The degree of a vertex $u \in V$ is denoted by $d_G(u)$, where each self-loop contributes 1 to the degree. For $S \subset V$, the *volume* of S is $\mathrm{vol}_G(S) = \sum_{u \in S} d_G(u)$. For $S \subset V$, the *cutset* of S, $C_G(S, \bar{S})$, is the multiset of edges with one endpoint in S and the other endpoint is $\bar{S}$. The *edge expansion* ρ_G and the *conductance* Φ_G of the graph G are

$$\rho_G = \min_{S \subset V, |S| \leq |V|/2} \frac{|C_G(S, \bar{S})|}{|S|}$$

and

$$\Phi_G = \min_{S \subset V, \mathrm{vol}_G(S) \leq \mathrm{vol}_G(V)/2} \frac{|C_G(S, \bar{S})|}{\mathrm{vol}_G(S)} \quad .$$

In Theorem 1 we establish constant conductance. Immediate implications for routing congestion and spectral gap are in Corollaries 3 and 4. The key technical ingredient in the proof of Theorem 1 is the bound of Lemma 2, which is time-invariant. This Lemma is established by a careful shifting argument that makes full use of the structure of the underlying evolutionary process.

Theorem 1 *For every positive constant integer $d \geq 2$, and for every positive constant $c < 2(d-1)-1$, there is a positive constant $\alpha = \alpha(d,c)$ such that the random graph $G_{d,n}$ has edge expansion α and conductance $\frac{\alpha}{d+\alpha}$, almost surely. In particular, for $\alpha < \min\{\frac{d-1}{2} - \frac{c+1}{4}, \frac{1}{5}, \frac{(d-1)\ln 2 - \frac{2}{5}\ln 5}{2(\ln d + \ln 2 + 1)}\}$,*

$$\Pr\left[\rho_{G_{d,n}} < \alpha\right] \leq o(n^{-c})$$

and

$$\Pr\left[\Phi_{G_{d,n}} \leq \frac{\alpha}{d+\alpha}\right] \leq o(n^{-c}) \quad .$$

Proof: Let us first bound conductance in terms of edge expansion. Let $S \subset [n]$ be a set with $\mathrm{vol}_{G_{d,n}}(S) \leq dn/2$. Since, by construction, every vertex associated with S contributes d to the total degree of S, we have $d|S| \leq \mathrm{vol}(S) \leq d|S| + C_{G_{d,n}}(S, \bar{S})$. The left hand side of this inequality implies $|S| \leq n/2$. Now the right hand side can be used to bound conductance by

$$\begin{aligned}
\Phi_{G_{d,n}} = & \min_{\substack{S \subset V \\ \mathrm{vol}_{G_{d,n}}(S) \leq dn/2}} \frac{C_{G_{d,n}}(S,\bar{S})}{\mathrm{vol}_{G_{d,n}}(S)} \\
\geq & \min_{\substack{S \subset V \\ \mathrm{vol}_{G_{d,n}}(S) \leq dn/2}} \frac{C_{G_{d,n}}(S,\bar{S})}{d|S| + C_{G_{d,n}}} \\
\geq & \frac{\rho}{d+\rho}
\end{aligned}$$

Now let us bound edge expansion. We will use a counting argument. Let us fix $k \leq n/2$ and let us fix a set $S \subset [n]$ with $|S| = k$. Let T_{dn} be the tree from which $G_{n,d}$ was generated. Say that a mini-vertex t, $1 \leq t \leq dn$ is GOOD if and only if either t is associated with S and the father of t is associated with $\bar{S}$, or t is associated with $\bar{S}$ and the father of t is associated with S. Say that a mini-vertex is BAD if and only if it is not GOOD. Realize that mini-vertex 1 is BAD, by convention. Realize also that if 1 belongs to S (resp. $\bar{S}$) then the first mini-vertex in $\bar{S}$ (resp. S) is always GOOD, by construction. Now let us fix the set $A \subset [dn]$ of GOOD mini-vertices, so that $|A| \leq \alpha k$. By Lemma 2

$$\Pr[\bigwedge_{\substack{t \in [dn] \\ t \notin A}} t \text{ is BAD}] \leq \frac{\binom{dk}{\alpha k}}{\binom{dn-\alpha k}{dk-\alpha k}} \quad .$$

There are $\binom{n}{k}$ choices for S. Once S is fixed, there are at most $\alpha k \binom{dn}{\alpha k}$ choices for A. Finally, because of the way we construct the graph we do not need to argue about singletons, therefore we need to

consider $2 \le k \le n/2$. The above imply

$$\begin{aligned}
\Pr[\rho_{G_{d,n}} < \alpha] &\le \textstyle\sum_{k=2}^{n/2} \binom{n}{k} \alpha k \binom{dn}{\alpha k} \frac{\binom{dk}{\alpha k}}{\binom{dn-\alpha k}{dk-\alpha k}} \\
&\le \textstyle\sum_{k=2}^{n/2} \alpha k \binom{dn}{\alpha k} \binom{dk}{\alpha k} \binom{(d-1)n-\alpha k}{(d-1)k-\alpha k}^{-1} && \text{, using } \binom{n}{k}\binom{(d-1)n-\alpha k}{(d-1)k-\alpha k} \le \binom{dn-\alpha k}{dk-\alpha k} \\
&\le \textstyle\sum_{k=2}^{n/2} \alpha k \left(\frac{n}{k}\right)^{\alpha k} \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{(d-1)k-\alpha k}{(d-1)n-\alpha k}\right)^{(d-1)k-\alpha k} && \text{, using the bound } \left(\frac{n}{k}\right)^k \le \binom{n}{k} \le \left(\frac{en}{k}\right)^k \\
&\le \textstyle\sum_{k=2}^{n/2} \alpha k \left(\frac{n}{k}\right)^{\alpha k} \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{k}{n}\right)^{(d-1)k-\alpha k} \\
&\le \textstyle\sum_{k=2}^{n/2} \alpha k \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{k}{n}\right)^{(d-1-2\alpha)k} .
\end{aligned}$$

There are $O(n)$ terms in the above summation. So we can bound the sum by $o(n^{-c})$, if we bound the leading term by $o(n^{-(c+1)})$. We thus need to study the function $f(k) = \alpha k \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{k}{n}\right)^{(d-1-2\alpha)k}$, $2 \le k \le n/2$. We will argue that, for some x in the interval $2 \le x \le n/2$, the function $f(k)$ is monotonically decreasing for $2 \le k \le x$ and monotonically increasing for $x \le k \le n/2$. To study the behavior of $f(k)$ we start by calculating the first derivative $f'(k) = \alpha \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{k}{n}\right)^{(d-1-2\alpha)k} \left[1 + k\left(2\alpha(1+\ln\frac{d}{\alpha}) + (d-1-2\alpha)(\ln\frac{k}{n}+1)\right)\right]$. Thus the sign of $f'(k)$ is determined by the function $g(k) = 1 + k\left(2\alpha(1+\ln\frac{d}{\alpha}) + (d-1-2\alpha)\left(\ln\frac{k}{n}+1\right)\right)$. We may verify that $g(2) < 0$ and $g(n/2) > 0$, for large enough n, while the first derivative of $g(k)$ is $g'(k) = 2\alpha(1+\ln\frac{d}{\alpha}) + (d-1-2\alpha)(\ln\frac{k}{n}+2)$ and the second derivative of $g(k)$ is $g''(k) = \frac{d-1-2\alpha}{k} > 0$. Hence the function $g(k)$ is convex and has exactly two roots, exactly one of which is in the interval $2 \le k \le n/2$. Consequently, for some x in the interval $2 \le x \le n/2$, the function $f(k)$ is monotonically decreasing for $2 \le k \le x$ and monotonically increasing for $x \le k \le n/2$. Thus, the leading term of $f(k)$, $2 \le k \le n/2$, can be bounded by the maximum of $f(2)$ and $f(n/2)$. We thus now require that $f(2)$ and $f(n/2)$ are both $o(n^{-(c+1)})$. First it can be verified that $f(2) = 2\alpha \left(\frac{ed}{\alpha}\right)^{4\alpha} \left(\frac{2}{n}\right)^{2(d-1-2\alpha)} = o(n^{-(c+1)})$ for $\alpha < \frac{d-1}{2} - \frac{c+1}{4}$ (it is easy to see that α is positive for the chosen ranges of d and c). Finally it can be also verified that $f(\frac{n}{2}) = \alpha\frac{n}{2}\left((\frac{ed}{\alpha})^{2\alpha}(\frac{1}{2})^{d-1-2\alpha}\right)^{\frac{n}{2}}$ which drops exponentially fast, for large enough n, as long as $(ed)^{2\alpha}\alpha^{-2\alpha} < 2^{d-1-2\alpha}$. For $\alpha < \frac{1}{5}$ the term $\alpha^{-2\alpha}$ is always smaller than $5^{\frac{2}{5}}$, and hence it suffices to bound $(ed)^{2\alpha}5^{\frac{2}{5}} < 2^{d-1-2\alpha}$. It can be verified that this is satisfied for $\alpha < \frac{(d-1)\ln 2 - \frac{2}{5}\ln 5}{2(\ln d + \ln 2 + 1)}$ (it is also easy to see that the numerator is always positive.) This completes the proof of Theorem 1. □

Lemma 2 *For a fixed subset $S \subset [n]$, $|S| = k$, and for a fixed subset $A \subset [dn]$, $|A| \le \alpha k$, the probability that all mini-vertices associated with $[dn] \setminus A$ are* BAD *in $G_{d,n}$ is at most $\binom{dk}{\alpha k}/\binom{dn-\alpha k}{dk-\alpha k}$.*

Proof: Let A_1 be the mini-vertices in A associated with S and A_2 be the mini-vertices in A

associated with $\bar{S}$. Let $|A_1|=k_1$ and $|A_2|=k_2$, with

$$k_1+k_2=|A| \ . \tag{1}$$

Let $x_1 < x_2 < \ldots < x_{dk-k_1}$ be the mini-vertices associated with S that do not belong to A. We may write $x_i = y_i + z_i + 1$, where y_i is the total number of mini-vertices that arrived prior to x_i and belong to A and z_i is the total number of mini-vertices that arrived prior to x_i and belong to $[dn] \setminus A$. Let $\bar{x}_1 < \bar{x}_2 < \ldots < \bar{x}_{dn-dk-k_2}$ be the mini-vertices associated with $\bar{S}$ that do not belong to A. We may write $\bar{x}_i = \bar{y}_i + \bar{z}_i + 1$, where $\bar{y}_i$ is the total number of mini-vertices that arrived prior to $\bar{x}_i$ and belong to A and $\bar{z}_i$ is the total number of mini-vertices that arrived prior to $\bar{x}_i$ and belong to $[dn] \setminus A$.

Now let us assume that the only GOOD mini-vertices are the ones belonging to A. Thus all mini-vertices associated with $[dn] \setminus A$ are BAD, and hence $x_1, \ldots, x_{dk-k1}$ as well as $\bar{x}_1, \ldots, \bar{x}_{dn-dk-k_2}$ are BAD. Recall also that the first mini-vertex 1 is not associated with A, since, by definition, 1 is BAD. Now realize that $\bigcup_{i=1}^{dk-k_1}\{z_i\} \cup \bigcup_{i=1}^{dn-dk-k_2}\{\bar{z}_i\} = \{0, 1, \ldots, dn-|A|-1\}$, and, equivalently,

$$\bigcup_{i=1}^{dk-k_1}\{z_i+1\} \cup \bigcup_{i=1}^{dn-dk-k_2}\{\bar{z}_i+1\} = [dn-|A|]. \tag{2}$$

Let us now proceed to bound the probability that all mini-vertices associated with $[dn] \setminus A$ are BAD, given that all mini-vertices in A are GOOD. First realize that the total volume of the graph when mini-vertex t arrives is $2(t-1)-1$, for $t \geq 2$. If $t = x_i$ (resp. $t = \bar{x}_i$), we can write this as

$$2(z_i + y_i) - 1 \qquad \text{resp.} \quad 2(\bar{z}_i + \bar{y}_i) - 1 \ . \tag{3}$$

We shall bound the probability that a mini-vertex in $[dn] \setminus A \setminus S$ is BAD and a mini-vertex in $[dn] \setminus A \setminus \bar{S}$ is BAD separately. Assume, without loss of generality, that $1 \notin S$, otherwise rename S and $\bar{S}$ (this lemma does not require $k < n/2$). It now follows that, by connectivity, the first mini-vertex in S is necessarily GOOD and thus this mini-vertex belongs to A. When x_i arrives, the total volume of S is contributed by: (a)All BAD mini-vertices that arrived prior to x_i and are associated with S, where each such mini-vertex contributes degree 2 and there are $i-1$ such mini-vertices. (b)All GOOD mini-vertices that arrived prior to x_i, where each such mini-vertex contributes 1 to the degree and there are y_i such mini-vertices; notice $y_i \geq 1$ since we argued above that the first mini-vertex in S belongs to A. This gives that the total degree of S when x_i arrives is

$$2(i-1) + y_i \tag{4}$$

Now (3) and (4) bound the probability that x_i attaches to S and is hence BAD, given that all mini-verties that arrived prior to x_i and belong to A are GOOD, while those belonging to $\bar{A}$ are BAD

by

$$\begin{aligned}
\tfrac{2(i-1)+y_i}{2(z_i+y_i)-1} &\leq \tfrac{2(i-1)+y_i}{2z_i+y_i} && \text{by subtracting } y_i-1 \geq 0 \text{ from the denomenator,} \\
& && \text{which is possible since } y_i \geq 1, \\
&\leq \tfrac{2(i-1)+|A|}{2z_i+|A|} && \text{by adding } |A|-y_i \geq 0 \\
& && \text{to the numerator and the denomenator,} \\
& && \text{which is possible since } y_i \leq |A|, \\
&\leq \tfrac{2i+|A|}{2z_i+2+|A|} && \text{by adding 2 to the numerator} \\
& && \text{and the denomenator,} \\
&= \tfrac{i+|A|/2}{z_i+1+|A|/2} \\
&= \tfrac{i+|A|}{z_i+1+|A|} && \text{by adding } |A|/2 \text{ to the numerator} \\
& && \text{and the denomenator.}
\end{aligned} \tag{5}$$

When $\bar{x}_i$ arrives, the total volume of $\bar{S}$ is contributed by: (a)All BAD mini-vertices that arrived prior to $\bar{x}_i$ and are associated with $\bar{S}$, where there are $i-1$ such mini-vertices and each one contributes degree 2 to $\bar{S}$, except for mini-vertex 1 which contributes degree 1. (b)All GOOD mini-vertices that arrived prior to $\bar{x}_i$, where each such mini-vertex contributes 1 to the degree and there are $\bar{y}_i$ such mini-vertices. This gives that the total degree of $\bar{S}$ when $\bar{x}_i$ arrives is

$$2(i-1)-1+\bar{y}_i \tag{6}$$

Now (3) and (6) bound the probability that $\bar{x}_i$ attaches to $\bar{S}$ and is hence BAD, given that all mini-verties that arrived prior to $\bar{x}_i$ and belong to A are GOOD, while those belonging to $\bar{A}$ are BAD by

$$\begin{aligned}
\tfrac{2(i-1)-1+\bar{y}_i}{2(\bar{z}_i+\bar{y}_i)-1} &\leq \tfrac{2(i-1)-1+\bar{y}_i}{2\bar{z}_i+\bar{y}_i-1} && \text{by subtracting } \bar{y}_i \geq 0 \text{ from the denomenator,} \\
&\leq \tfrac{2(i-1)-1+|A|}{2\bar{z}_i+|A|-1} && \text{by adding } |A|-\bar{y}_i \geq 0 \\
& && \text{to the numerator and the denomenator,} \\
& && \text{which is possible since } \bar{y}_i \leq |A|, \\
&\leq \tfrac{2i+|A|}{2\bar{z}_i+2+|A|} && \text{by adding 3 to the numerator} \\
& && \text{and the denomenator,} \\
&= \tfrac{i+|A|/2}{\bar{z}_i+1+|A|/2} \\
&= \tfrac{i+|A|}{\bar{z}_i+1+|A|} && \text{by adding } |A|/2 \text{ to the numerator} \\
& && \text{and the denomenator.}
\end{aligned} \tag{7}$$

Now (5) and (7) imply that the probability that all mini-vertices associated with $[dn] \setminus A$ are BAD,

given that all mini-vertices in A are GOOD, is at most

$$\begin{aligned}
\prod_{i=1}^{dk-k_1} \frac{i+|A|}{z_i+1+|A|} \quad \prod_{i=1}^{dn-dk-k_2} \frac{i+|A|}{\bar{z}_i+1+|A|} &= \frac{\prod_{i=1}^{dk-k_1}(i+|A|) \quad \prod_{i=1}^{dn-dk-k_2}(i+|A|)}{\prod_{i=1}^{dn-|A|}(i+|A|)} \quad , \text{ using (2)} \\
&= \frac{\prod_{i=1}^{dk-k_2+|A|} i \quad \prod_{i=1}^{dn-dk-k_2+|A|} i}{(dn)! \ \ |A|!} \quad , \quad \begin{array}{c}\text{by myltiplying} \\ \text{numerator and} \\ \text{denomenator} \\ \text{with } (|A|!)^2\end{array} \qquad (8) \\
&= \frac{(dk+k_2)! \ \ (dn-dk+k_1)!}{(dn)! \ \ |A|!} \quad , \text{ using (1)} \\
&= \prod_{i=0}^{k_2-1} \frac{dk+k_2-i}{dn-i} \prod_{i=0}^{k_1-1} \frac{dn-dk+k_1-i}{dn-k_2-i} \quad \frac{(dk)! \ \ (dn-dk)!}{(dn-|A|)! \ |A|!} \\
&\leq \frac{(dk)! \ \ (dn-dk)!}{(dn-|A|)! \ \ (|A|)!}
\end{aligned}$$

where the last inequality of (8) follows by (9) and (10) below.

$$\frac{dk+k_2-i}{dn-i} \leq 1 \qquad , \qquad 0 \leq i \leq k_2-1 \quad . \tag{9}$$

$$\frac{dn-dk+k_1-i}{dn-k_2-i} \leq 1 \qquad , \qquad 0 \leq i \leq k_1-1 \quad . \tag{10}$$

To see (9) it suffices to argue that $dk+k_2 \leq dn$, equivalently, that $k_2 \leq dn-dk$, which is true since k_2 is the cardinality of A_2 (mini-vertices associated with $\bar{S}$ that also belong to A) and $dn-dk$ is the cardinality of all the mini-vertices associated with $\bar{S}$, which is a set that includes A_2. To see (10) it suffices to argue that $dn-dk+k_1 \leq dn-k_2$, equivalently, that $k_1+k_2 \leq dk$, which is true since $k_1+k_2=|A|$ following (1), and $|A| \leq \alpha k \leq dk$ for the used range of α.

Finally, we can bound (8) as in the statement of Lemma 2 by noting:

$$\begin{aligned}
\frac{(dk)!\,(dn-dk)!}{(dn-|A|)!\,|A|!} &= \frac{(dk-|A|)!\,(dn-dk)!}{(dn-|A|)!} \cdot \frac{(dk)!}{(|A|)!\,(dk-|A|)!} \quad , \quad \begin{array}{c}\text{by multiplying} \\ \text{numerator and} \\ \text{denomenator} \\ \text{with } (dk-|A|!)!\end{array} \\
&= \binom{dk}{|A|}\binom{dn-|A|}{dk-|A|}^{-1} \\
&\leq \binom{dk}{\alpha k}\binom{dn-\alpha k}{dk-\alpha k}^{-1}.
\end{aligned}$$

This completes the proof of Lemma 2. □

We may now quote approximation techniques for multicommodity flow [33, 45] and claim:

Corollary 3 *Let $G_{d,n}$ be a random graph as in Theorem 1. There is a polynomial time algorithm that routes $d_{G_{d,n}}(u) \cdot d_{G_{d,n}}(v)$ units of flow between every pair of vertices u and v, with maximum link congestion $O(n \log n)$.*

The reason that we insist of $d_{G_{d,n}}(u) \cdot d_{G_{d,n}}(v)$ units of flow is that, in general (e.g. for large d), the random graph may model the core of the entire network. In that case, every node in the core has to serve a number of customers that tends to be proportional to its degree in the core, hence the demand between two nodes in the core becomes proportional to the product of their degrees (we refer the reader to [27] for further explanation of the assumptions on uniform demand and capacities, and the implications of Corollary 3 in routing congestion on the Internet).

Most of the routing on the Internet is done along integral shortest paths [28]. Leighton and Rao have already observed that randomized rounding applies to their algorithm, hence Corollary 3 can be restated for integral paths. We can also apply the techniques of disjoint paths for constant-degree expanders (e.g. [12]) and for routing along short paths [31] through the following simple construction: Every vertex u in $G_{d,n}$ of degree $d_{G_{d,n}}(u)$ is replaced with $d_{G_{d,n}}(u)$ mini-vertices. Each mini-vertex is connected to the corresponding edge of $G_{d,n}$, and within the $d_{G_{d,n}}(u)$ mini-vertices we put a constant-degree expander. It can be argued routinely that the resulting graph is a constant degree expander.

Another notable implication of Theorem 1 concerns the spectral gap of the stochastic normalization of the adjacency matrix of the graph[1]. In particular, since we know that $\lambda_2 < 1 - \frac{\Phi^2}{2}$ (e.g. see [43], page 53), we get:

Corollary 4 *Let $G_{d,n}$ be a random graph as in Theorem 1. Let A be the adjacency matrix of $G_{d,n}$. Let P be the stochastic matrix corresponding to a random walk in $G_{d,n}$. The largest eigenvalue of P is $\lambda_1 = 1$. Let λ_2 be the second largest eigenvalue. Then, for some positive constant c, the second eigenvalue $\lambda_2 < 1 - c$, almost surely.*

It is known that the cover time of a graph is bounded by $O(\frac{n \log n}{1-\lambda_2})$ —e.g. see [13]. Then Corollary 4 gives cover time $O(n \log n)$. We note that the cover time of scale free graphs has been associated with crawling and searching on the world-wide web and P2P networks [20, 19, 1].

Constant-degree expander graphs have played a central role in algorithms and complexity over the last thirty years [37, 43, 15, 45, 40]. In a rather strong sense, Theorem 1 and Corollary 4 suggest analogies between constant-degree constant expanders and constant average degree scale free graphs. It is reasonable to expect that analogies will find many further applications.

3 On the Frugality of Random Graphs

For any graph G and vertices s and t, consider the shortest path P from s to t (assumed to be unique, with ties broken lexicographically). For each edge e on this path we define the *Vickrey–Clarke–Groves (or VCG) overpayment* of e with respect to s and t, denoted $v(e, s, t)$, to be the increase in the length of the shortest path from s to t if edge e were deleted. If e is a bridge disconnecting s from t, $v(e, s, t) = 0$.

Our nonstandard way of dealing with bridges in our definition needs some explanation. It allows us to analyze with respect to this metric the standard random graph models in which bridges and small components are present with some probability even in reasonably dense graphs. Besides, our

[1] This is not to be confused with the spectrum of the adjacency matrix prior to stochastic normalization, considered elsewhere [24, 36, 18]. The eigenvalues of the matrix prior to normalization are a restatement of skewed statistics in the large degrees, and are hence of no particular content or algorithmic significance [36].

definition is compatible with the premises of the experimental result that we are seeking to explain: In [25] it was pointed out that, in the graph of the Internet's autonomous systems, $v(e,s,t)$ is between 0.3 and 0.4 on the average *if restricted to the biconnected core of the Internet graph.* (That experiment considered vertices as the costly elements, while we count the number of edges; our results are trivially translatable to the vertex case.)

It is reasonable to consider the VCG overpayment as a parameter that somehow reflects the degree to which a network is "monopolistic". (The bridges, which we ignore in our calculation, also exist in the Internet, as deliberate decisions of autonomous systems to depend on a single provider, and it is reasonable to consider them a phenomenon quite distinct from large imbalances in path lengths.) A variety of recent negative results, reviewed in the introduction, establish that biconnected graphs can have terribly high overpayments even in much more relaxed models than VCG. For the case of VCG, it is easy to see that the overpayment can be as high as a factor of k for a cycle of length k. We can show the following:

Theorem 5 *For vertices s and t in $G \in G_{n,p}$ (the Erdős-Renyi random graph model with n nodes and edge probability p) with $np = \omega(\log n)$, the average $v(e,s,t)$ of edges on a shortest path between s and t is $O(1)$.*

Let us first introduce some notation. Consider the sequence of breadth-first search frontiers (sets of unexpanded nodes) around s and t. Define $\Gamma_i(s) = \{v \in V(G) : \text{distance}(s,v) = i\}$ and $N_i(s) = \bigcup_{j=0}^{i} \Gamma_j(s)$. Define $\Gamma_i(t)$ and $N_i(t)$ similarly. Throughout this section, we will use the lemma (see [9], Lemma 10.7) that for every v and any $\delta > 0$, for n sufficiently large and i less than or equal to $d = \log n / \log np$ with probability $1 - O(n^{-8})$, we have

$$||\Gamma_i(v)| - (pn)^i| \leq \delta (pn)^i. \tag{11}$$

Let $\Pi \subset G_{n,p}$ be the set of graphs satisfying inequality (11) for $\delta = \frac{1}{8}$ for vertices s and t .

By inequality (11), if $G \in \Pi$, $|\Gamma_d(s)| + |\Gamma_d(t)| > n+2$ and therefore the breadth-first search frontiers $N_i(s)$ and $N_i(t)$ meet at least twice for some $i \leq d$. Let P be the set of the edges in the shortest path between s and t. $N_i(s)$ and $N_i(t)$ meet for the first time at $i = \lfloor (|P|+1)/2 \rfloor$. Let $P_2(s)$ be the path defined by our breadth first search connecting the second meeting point to s. Define $P_2(t)$ similarly and let $P_2 = P_2(s) \cup P_2(t)$.

Lemma 6 *For $G \in \Pi$, $E(|P_2| - |P|)$ is $O(1)$.*

Proof: Let k be the smallest i such that $|\Gamma_i(s)| \geq \sqrt{n}$ (where $\sqrt{n}$ comes from the birthday paradox). We will bound the expected value of $|P_2| - 2k$ and $2k - |P|$ by $O(1)$, separately and then will have $E(|P_2| - |P|) = E(|P_2| - 2k) + E(2k - |P|)$.

For P_2, consider the set $\Gamma_{k+i}(s)$, and $\Gamma_{k+i}(t)$, $i = 0, 1, \ldots, d-k-1$, and calculate the probability that $\Gamma_{k+i}(s)$ and $\Gamma_{k+i}(t)$ intersect in fewer than two points. For graphs in Π, by inequality (11), these sets have at least $g = \frac{3}{4}\sqrt{n}(pn)^i$ elements, and hence, in these graphs, the probability that the sets intersect in fewer than two points is at most

$$(1 + \frac{g^2}{n})(1 - \frac{g-1}{n})^{g-1} \leq 2(pn)^{2i} exp\{-\frac{9}{16}(pn)^{2i}\}.$$

Thus, the expectation:

$$E(|P_2|-2k) = \sum_i \Pr[|P_2| > i+k] \geq \sum_i \Pr[|\Gamma_{i+k}(s) \cap \Gamma_{i+k}(t)| < 2] \sum_{i>0} 2(pn)^{2i} exp\{-\frac{9}{16}(pn)^{2i}\} = O(1)$$

Now, in order to bound the expectation of $2k - |P|$, consider the sets $\Gamma_{k-i}(s), \Gamma_{k-i}(t)$, $i = 1, \ldots, k$. We will bound the probability that $\Gamma_{k-i}(s)$ and $\Gamma_{k-i}(t)$ meet before time k. Again, by (11), the cardinality of these sets is at most $h_i = \frac{5}{4}\sqrt{n}/(np)^{i-1}$. The probability that they intersect is at most

$$1 - (1 - \frac{h_i}{n})^{h_i} \leq \min\{1, 1 - (1 - (\frac{h^2}{n}))\} \leq \min\{1, 25 \ over 16(np)^{2-2i}\}.$$

Hence, the expectation of $k - |P|$ is at most $4\sum_{i>0} \min\{1, (np)^{2-2i}\} = O(1)$, which completes the proof of the upper bound. □

Lemma 7 *For P and P_2 defined above in $G \in \Pi$, $E(|P_2 \cap P|) = O(1)$.*

Proof: It is easy to see that $|P \cap P_2(s)| \geq i$ implies that P and P_2 go through the same vertex of $\Gamma_j(s)$ for some $j \geq i$. But since the vertices of $\Gamma_j(s)$ through which either P or P_2 pass can be viewed as independent uniform draws, the probability that they are the same in $\Gamma_j(s)$ is precisely $1/|\Gamma_j(s)|$. Therefore, $E(|P_2(s) \cap P|) \leq E(\sum_{i\geq 1} 1/|\Gamma_i(s)|)$. Using the same argument for $P_2(t)$ and inequality (11) we have

$$E(|P_2 \cap P|) \leq E(\sum_{i\geq 1} \frac{2}{|\Gamma_i(t)|}) = O(1).$$

□

Proof of Theorem 5 : Since $v(e, s, t)$ is always bounded by n, we can restrict our calculation to a subset of graphs with measure $1 - o(1/n)$. So we can assume that $G \in \Pi$.

For all edges in $P \setminus P_2$, $v(e, s, t) \leq |P_2| - |P|$. For an edge $e \in P_2 \cap P$, $v(e, s, t) \leq \text{diameter}(G \setminus e)$. So we have

$$|P_2| - |P| \leq \frac{\sum_{e\in P} v(e, s, t)}{|P|} \leq |P_2| - |P| + \text{diameter}(G \setminus e)\frac{|P \cap P_2|}{|P|} \tag{12}$$

both $|P|$ and $\text{diameter}(G \setminus e)$ are $\Theta(\log n/\log np)$ with probability $1 - o(1/n)$ (for example see [9]). We also know that for $G \in \Pi$ both $E(|P \cap P_2|)$ and $E(|P_2| - |P|)$ are O(1) by Lemmas 7 and 6. □

Extensions

It would be of interest to extend this result to more "Internet-like" random graph models whose degree sequences follow power-law distributions, such as the model of growth with preferential attachment discussed in the previous section or power-law random graphs [2]. The difficulty with such graphs arises because the sizes of the sets $\Gamma_i(v)$ for a vertex v do not behave as nicely as in the inequality (11); in fact the variances can be very large. It would also be of interest to extend this result to very sparse graphs such as Erdős-Renyi $G(n, p)$ for constant expected degree and for random constant-degree regular graphs; we believe that the same techniques and similar calculations as in Theorem 5 will suffice for this case.

Acknowledgments

We wish to thank Miki Ajtai and Elchanan Mossel for useful discussions. The first and last authors were supported by NSF undr ITR-0220343. The second author was supported by NSF under grant CCR-01215555.

References

[1] Adamic, L., Lukose, R., Puniyani, A. and Huberman, B., "Search in Power Law Networks", Physical review E, Volume 64.

[2] Aiello, W., Chung, F.R.K. and Lu, L., "A Random Graph Model for Power Law Graphs", FOCS 2000, pp. 171-180.

[3] Aiello, W., Chung, F.R.K. and Lu, L., "Random Evolution in Massive Graphs", FOCS 2001, pp 510-519.

[4] Albert, R. and Barabási, A.-L., "Error and Attack Tolerance of Complex Networks", Nature 406, 2000.

[5] Archer, A. and Tárdos, E. "Frugal Path Mechanisms", SODA 2002, pp 991-999.

[6] Barabási, A.-L. and Albert, R., "Emergence of scaling in random graphs", Science 286 (1999), pp. 509-512.

[7] Bar-Yossef, Z., Berg, A., Chien, S., Fakcharoenphol, J. and Weitz, D., "Approximating Aggregate Queries about Web Pages via Random Walks", VLDB 2000.

[8] Bollobás, B. and de la Vega, F. W., "The Diameter of Random Regular Graphs", Combinatorica 2(2): 125-134 (1982).

[9] Bollobás, B., "Random Graphs", Cambridge University Press, 2nd edition (September 15, 2001).

[10] Bollobás and Riordan, O., "The diameter of scale-free graphs", to appear in Combinatorica.

[11] Bollobás, B., Riordan, O., Spencer, J. and Tusnády, G., "The degree sequence of a scale-free random graph process", Random Structures and Algorithms,Volume 18, Issue 3, 2001, pp 279-290.

[12] Broder, A., Frieze, A. and Upfal, E., SIAM Journal of Computing, (23) No 5, pp 976-989.

[13] Broder, A. and Karlin, A., "Bounds on the Cover Time", J. Theoretical Probability 2(1) (1989), pp 101-120.

[14] Chen, Q., Chang, H., Govindan, R., Jamin, S., Shenker, S. and Willinger, W., "The Origins of Power-Laws in Internet Topologies Revisited", Infocomm 2002.

[15] Chung, F.R.K., "Spectral Graph Theory", AMS-RCSM, No 92, 1997.

[16] Chung, F.R.K. and Lu, L., "The Diameter of Random Sparse Graphs", Advances in Applied Math., 26 (2001), 257-279.

[17] Chung, F.R.K. and Lu, L., "The average distance in a random graph with given expected degrees", Proceedings of the National Academy of Sciences, 99, No 25, December 2002.

[18] Chung, F.R.K., Lu, L. and Vu. V., "The spectra of random graphs with given expected degrees", Proceedings of the National Academy of Sciences (to appear).

[19] Cooper, C. and Frieze, A., "A general model for web graphs", *Proceedings of ESA*, 2001, pp 500-511.

[20] Cooper, C. and Frieze, A., "Crawling on Web Graphs", STOC 2002, pp 419-427.

[21] Cooper, C. and Frieze, A., "The cover time of sparse random graphs", SODA 2003, pp 148-157.

[22] Elkind, E. and Sahai, A. and Steiglitz, K. "Frugality in Path Auctions", Manuscript, 2003.

[23] Fabrikant, A., Koutsoupias, E. and Papadimitriou, C.H. "Heuristically Optimized Tradeoffs", ICALP 2002.

[24] Faloutsos, M., Faloutsos, P. and Faloutsos, C., "On Power-law Relationships of the Internet Topology", In Proceedings *Sigcomm* 1999, pp 251-262.

[25] Feigenbaum, J., Papadimitriou, C.H., Sami, R. and Shenker, S., "A BGP-based Mechanism for Lowest-Cost Routing"," Proceedings of the 21st Symposium on Principles of Distributed Computing, New York, 2002, pp. 173-182.

[26] Gkantsidis, C., Mihail, M. and Zegura, E., "Spectral Analysis of Internet Topologies", INFOCOM 2003 (to appear).

[27] Gkantsidis, C., Mihail, M. and Saberi, A., "Throughput and Congestion in Power Law Graphs", Sigmetrics 2003 (to appear).

[28] Griffin, T., "An Introduction to Interdomain Routing and BGP", SIGCOMM 2001 Tutorial.

[29] Holme, P., "Edge Overload Breakdown in Evolving Networks", Physical Review E 66, 2002.

[30] Holme, P. amd Kim, B.J., "Attack Vulnerability of Complex Networks", Physical Review E 65, 2002.

[31] Kleinberg, J. and Rubinfeld, R., "Short Paths in Expander Graphs", FOCS 96, pp 86-95.

[32] Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A. and Upfal, E., "Stochastic models for the Web graph", FOCS 2000, pp 57-65.

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

[34] Public Route Server and Looking Glass Site List, Traceroute.org, http://www.traceroute.org.

[35] Mas-Collel, A., Whinston, M. and Green, J., "Microeconomics Theory", Oxford University Press, 1995.

[36] Mihail, M. and Papadimitriou, C.H., "On the Eigenvalue Power Law", RANDOM 2002.

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

[38] National Laboratory for Applied Networking Research, Routing Data, http://moat.lanr.net/Routing/rawdata.

[39] Nisan, N. and A. Ronen, "Algorithmic Mechanism Design" STOC 1999, pp 129–140.

[40] Papadimitriou, C.H., "Computational Complexity", Addison-Wesley, 1994.

[41] Papadimitriou, C.H., "Algorithms, Games, and the Internet", Keynote Address, STOC 2001, pp 749-753.

[42] Papadimitriou, C.H., "Game Theory and Math Economics: A TCS Introduction", Tutorial, FOCS 2001.

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

[44] Towsley, D., "Modeling the Internet: Seeing the Forest through the Trees", Keynote Address, Sigmetrics 2002.

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

[46] Willinger, W., Govindan, R., Paxson, V. and Shenker, S., "Scaling Phenomena in the Internet: Critically examining Criticality", Proceedings of the National Academy of Sciences, Vol 99, 2002.