

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

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 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, 30]. 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 fac-

tor 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 2.2). 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 2.1) 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 an 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 and below by a non-increasing function of the expected degree (we conjecture that it grows as the inverse of the expected degree). We show this in the $G_{n,p}$ model, first when $np = \omega(\log n)$, by a careful analysis of the shortest and second shortest paths in a random graph. We then extend it down to constant np by a technique that bounds the size of “long” breadth-first neighborhoods in random graphs. These results can be extended routinely to the d -regular model. As for the “scale-free” models, using results from [17] it is easy to show that the frugality of random graphs with specified power law degree sequences have $O(\log \log n)$ VCG overpayment. We conjecture that constant upper and lower bounds hold in suitably defined scale-free models.

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 probabil-

ity 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 $\text{vol}_G(S) = \sum_{u \in S} d_G(u)$. 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

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

In Theorem 2.1 we establish constant conductance. Immediate implications for routing congestion and spectral gap are in Corollaries 2.3 and 2.4. The key technical ingredient in the proof of Theorem 2.1 is the bound of Lemma 2.2, 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 2.1 *There is a positive constant α such that, for any constant $d \geq 2$, the random graph $G_{d,n}$ has edge expansion α and conductance $\frac{\alpha}{d+\alpha}$, almost surely. In particular, for any non-negative constant $c < 2(d-1) - 4\alpha - 1$*

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

and

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

Proof: Let us first bound conductance in terms of edge expansion. Let $S \subset [n]$ be a set with $\text{vol}_{G_{d,n}}(S) \leq dn/2$. Since, by construction, every vertex associated with S contributes d to the total degree of S , we have $d|S| \leq$

$\text{vol}(S) \leq d|S| + C_{G_{d,n}}(S, \bar{S})$. The left hand side of this inequality implies $|S| \leq n/2$. Now the right hand side can be used to bound conductance by

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

Now let us bound edge expansion. We will use a counting argument. Say that a set $S \subset [n]$ is BAD if $|C_{G_{d,n}}(S, \bar{S})| < \alpha|S|$. Now 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 $G_{d,n}$ was generated. Say that a mini-vertex t , $1 \leq t \leq dn$ is BAD 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 GOOD if and only if it is not BAD. 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 2.2

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

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 $2 \leq k \leq n/2$. The above imply

$$\Pr[\exists \text{ a BAD SET } S] \leq \sum_{k=2}^{n/2} \binom{n}{k} \alpha k \binom{dn}{\alpha k} \frac{\binom{dk}{\alpha k}}{\binom{dn-\alpha k}{dk-\alpha k}}$$

using $\binom{n}{k} \binom{(d-1)n-\alpha k}{(d-1)k-\alpha k} \leq \binom{dn-\alpha k}{dk-\alpha k}$,

$$\begin{aligned} &\leq \sum_{k=2}^{n/2} \alpha k \binom{dn}{\alpha k} \binom{dk}{\alpha k} \left(\frac{(d-1)n-\alpha k}{(d-1)k-\alpha k}\right)^{-1} \\ &\leq \sum_{k=2}^{n/2} \alpha k \left(\frac{n}{k}\right)^{\alpha k} \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{(d-1)k-\alpha k}{(d-1)n-\alpha k}\right)^{(d-1)k-\alpha k} \end{aligned}$$

using the bound $\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k$,

$$\begin{aligned} &\leq \sum_{k=2}^{n/2} \alpha k \left(\frac{n}{k}\right)^{\alpha k} \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{k}{n}\right)^{(d-1)k-\alpha k} \\ &\leq \sum_{k=2}^{n/2} \alpha k \left(\frac{ed}{\alpha}\right)^{2\alpha k} \left(\frac{k}{n}\right)^{(d-1)k-2\alpha k}. \end{aligned}$$

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 $o(n^{-(c+1)})$. It can be seen that, for α small enough, all terms are smaller than the term for $k = 2$, provided $d-1-2\alpha > 0$, which is true for α small enough and $d \geq 2$. Hence we need to bound $n^{-(2d-1-4\alpha)}$. This can be bounded by $n^{-(c+1)}$, for c as in the statement of the theorem. \square

Lemma 2.2 *For a fixed subset $S \subset [n]$, $|S| = k$, and for a fixed subset $A \subset [dn]$, $|A| \leq \alpha k$, 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 $k_1 + k_2 = |A|$. Let $x_1 < x_2 < \dots < x_{dk-k_1}$ 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 $\bar{x}_1 < \bar{x}_2 < \dots < \bar{x}_{dk-k_2}$ 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 $[dn] \setminus L$. The crucial observation is that

$$\bigcup_{i=1}^{dk-k_1} \{z_i\} \cup \bigcup_{i=1}^{dn-dk-k_2} \{\bar{z}_i\} = [dn-dk] \quad (1)$$

We now bound the probability that all mini-vertices associated with $[dn] \setminus A$ are BAD, given that all mini-vertices in A are GOOD. Note that the first mini-vertex 1 is not associated with A , since, by definition, 1 is BAD. Assume, without loss of generality, that $1 \notin S$, otherwise rename S and \bar{S} (this lemma does not require $k < n/2$). Hence, by connectivity, the first mini-vertex in S is necessarily GOOD and we may assume that that mini-vertex belongs to A . Realize that the total volume of the graph when mini-vertex t arrives is

$$2(t-1) - 1, \quad \text{for } t \geq 2.$$

If $t = x_i$ (resp. $t = \bar{x}_i$), we can write this as

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

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

$$2(i-1) + y_i \quad (3)$$

Now (2) and (3) bound the probability that x_i attaches to S and is hence BAD, given that all mini-vertices that arrived prior to x_i and belong to A are GOOD, while those belonging to \bar{A} are BAD by

$$\begin{aligned} \frac{2(i-1)+y_i}{2(z_i-1)+2y_i-1} &\leq \frac{2(i-1)+y_i}{2(z_i-1)+y_i} \quad \text{since } y_i \geq 1 \\ &\leq \frac{2(i-1)+|A|}{2(z_i-1)+|A|} \\ &= \frac{(i-1)+|A|/2}{(z_i-1)+|A|/2} \\ &= \frac{i+|A|}{z_i+|A|} \end{aligned} \quad (4)$$

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

$$2(i-1) + \bar{y}_i \quad (5)$$

Now (2) and (5) bound the probability that \bar{x}_i attaches to \bar{S} and is hence BAD, given that all mini-vertices that arrived prior to \bar{x}_i and belong to A are GOOD, while those belonging to \bar{A} are BAD by

$$\begin{aligned} \frac{2(i-1)-1+\bar{y}_i}{2(\bar{z}_i-1)-1+2\bar{y}_i} &\leq \frac{2(i-1)-1+\bar{y}_i}{2(\bar{z}_i-1)-1+\bar{y}_i} \\ &\leq \frac{2(i-1)-1+|A|}{2(\bar{z}_i-1)-1+|A|} \quad (6) \\ &\leq \frac{i+|A|}{\bar{z}_i+|A|} \end{aligned}$$

Now (4) and (6) imply that the probability that all mini-vertices not belonging to S are BAD is at most

$$\prod_{i=1}^{dk-|A|} \frac{i+|A|}{z_i+|A|} \prod_{i=1}^{dn-dn-|A|} \frac{i+|A|}{z_i+|A|}$$

Which is clearly

$$\begin{aligned} &= \frac{(dk)!(dn-dk)!}{(dn-|A|)! (|A|)!} \quad \text{using (1)} \\ &= \frac{(dk-|A|)!(dn-dk)!}{(dn-|A|)!} \cdot \frac{(dk)!}{(|A|)!(dk-|A|)!} \\ &= \binom{dk}{|A|} \binom{dn-|A|}{dk-|A|}^{-1} \\ &\leq \binom{dk}{\alpha k} \binom{dn-\alpha k}{dk-\alpha k}^{-1}. \end{aligned}$$

□

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

Corollary 2.3 *Let $G_{d,n}$ be a random graph as in Theorem 2.1. 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 2.3 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 2.3 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 2.1 concerns the spectral gap of the stochastic normalization of

the adjacency matrix of the graph¹. In particular, since we know that $\lambda_2 < 1 - \frac{\Phi^2}{2}$ (e.g. see [43], page 53), we get:

Corollary 2.4 *Let $G_{d,n}$ be a random graph as in Theorem 2.1. 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 2.4 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 2.1 and Corollary 2.4 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 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*

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].

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 3.1 *For $G \in G_{n,p}$ (the Erdős-Renyi random graph model with n nodes and edge probability p) with $np = \omega(\log n)$, with probability $O(n^{-c})$ for some $c > 0$, the average $v(e, s, t)$ over all vertices s, t and edge e on the shortest path between s and t is $O(1)$ and $\Omega(\frac{1}{np})$.*

Proof: Let P be the set of edges in the shortest path and P_2 be the set of edges of any other path connecting s and t , i.e. $P \neq P_2$. 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

$$\frac{\sum_{e \in P} v(e, s, t)}{|P|} \leq |P_2| - |P| + \text{diameter}(G \setminus e) \frac{|P \cap P_2|}{|P|} \quad (7)$$

Since both $|P|$ and $\text{diameter}(G \setminus e)$ are $\Theta(\log n / \log np)$ with probability $1 - o(1/n)$, we must find a P_2 for which $|P \cap P_2|$ is bounded from above by a constant in expectation.

Consider the sequence of breadth-first search frontiers (sets of unexpanded nodes) $\Gamma_i(s) = \{v \in V(G) : \text{distance}(v, s) = i\}$ and $N_i(s) = \bigcup_{j=0}^i \Gamma_j(s)$. By Lemma ([9], Lemma 10.7), for every s and any $\delta > 0$, with probability $1 - O(n^{-8})$ we have

$$||\Gamma_i(s)| - (pn)^i| \leq \delta(pn)^i. \quad (8)$$

We can define $\Gamma_i(t)$ and $N_i(t)$ similarly. The frontiers $\Gamma_i(s)$ and $\Gamma_i(t)$ meet for the first time at $i = \lfloor (|P| + 1)/2 \rfloor$ (ignoring an additive constant of 1, for a moment). We will continue growing these frontiers until they meet again. Let $P_2(s)$ be the path defined by our

breadth first search connecting this second meeting point to s . Define $P_2(t)$ similarly and let $P_2 = P_2(s) \cup P_2(t)$.

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 equation (8) we have with high probability

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

In order to bound $|P_2| - |P|$, let k be the smallest i such that $|\Gamma_i(s)| \geq \sqrt{n}$ (where \sqrt{n} comes from the birthday paradox). It suffices to consider the case in which $|P_2| \geq 2k$ and $|P| \leq 2k$, and calculate and add the expectations of the deviations from these bounds. For P_2 , consider the set $\Gamma_{k+i}(s)$, and $\Gamma_{k+i}(t)$, $i = 0, 1, \dots, n$, and calculate the probability that they intersect in fewer than two points. For all but a fraction of $O(n^{-8})$ of all graphs, 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 $g(1 - g/n)^{g-1} \leq e^{-(pn)^{2i/3}}$. Thus, the expectation of $|P_2| - 2k$ is at most $2 \sum_{i > 0} e^{-(pn)^{2i/3}} = O(1)$.

Now, in order to bound the expectation of $2k - |P|$, consider the sets $\Gamma_{k-i}(s), \Gamma_{k-i}(t)$, $i = 1, \dots, k$, of cardinality at most $h = \frac{5}{4}\sqrt{n}/(np)^{i-1}$. The probability that these sets intersect is at most $1 - (1 - h/n)^h \leq 1 - e^{-1/2(np)^{2i-2}} \leq \max\{1, (np)^{2-2i}\}$. Hence, the expectation of $k - |P|$ is at most $2 \sum_{i > 0} \max\{1, (np)^{2-2i}\} = O(1)$, which completes the proof of the upper bound.

For the lower bound, it suffices to notice that the expectation of $k - |P|$ is $\Omega(1/np)$, and that indeed $|P| < k$ and $|P_2| > k$ with some non-vanishing probability. \square

We believe that with more careful calculations the upper bound can be made asymptotically equal to the present lower bound, and with constants for both that are quite small and close to each other. The inaccuracies of our present calculation stem from (a) the slight uncertainty in the size of Γ_k and (b) round-off problems in defining k , and appear to be susceptible to a better analysis.

Extensions

We would like to extend this result to less restrictive and more “Internet-like” random graph models. By routine techniques (see, for example [8]), a similar result can be shown for the random d -regular graph model.

We next focus on its extension to the case of *constant* expected degree. One first obstacle is that there is no known equivalent of [9]’s Lemma 10.7 in this case, since for constant expected degree one cannot use Chernoff bounds at each step of the breadth-first search. In fact, the breadth-first search tree can be skinny with quite high probability. The following Lemma is useful in this direction (and possibly in others, see [16] for a related result).

Lemma 3.2 *There is a constant $D \leq 96 \ln 2$ such that for any $1/2 \geq \epsilon > 0$ and $p \geq D/n$ for all but a fraction of $O(n^{-a})$ for some $a > 0$ of graphs in G_{np} and vertices s there is a $k \leq h\epsilon \log n / \log(np/2)$ for some $h > 2$ such that, conditioned on $\Gamma_k \neq \emptyset$, $|\Gamma_k| \geq n^\epsilon$.*

Proof: Assume without loss of generality that $np = D$ and $k = h\epsilon \log n / \log d$ where $d = D/2$. Consider the sequence $\Gamma_0, \Gamma_1, \dots, \Gamma_k$ and call $i \leq k$ a *small step* if $|\Gamma_i|/|\Gamma_{i-1}| \leq d$.

We have:

$$P[|\Gamma_k| \leq n^\epsilon] \leq$$

(since, in order for this to happen, and since no Γ_i is empty, there must be a small step in the last $\epsilon \log n / \log d$ steps)

$$P[\text{there is a small step in the last } \epsilon \log n / \log d \text{ steps}] \leq$$

(by breaking the steps before the last small one into runs of $m > 0$ non-small steps of lengths L_i adding to at least $k' = k(h-1)/h$, and then taking Chernoff bounds of the form $P[|x - \mu| > \alpha\mu] < e^{-\alpha^2\mu/12}$ on the small steps)

$$\sum_{m=1}^k \binom{k}{m} \prod_{i=1}^m \exp\left\{-\left(D - d/D\right)^2 \frac{1}{12} d^{L_i}\right\} \leq$$

(by observing that each product is maximized by letting all L_i ’s be equal to k'/m)

$$\sum_{m=1}^{k'} \binom{k}{m} \exp\left\{-\frac{m}{48} d^{k'/m}\right\} \leq 2^k e^{-kD/96}$$

which is small for the values of D and k given in the statement. \square

Notice that our rough calculation above only guarantees exponentially growing breadth-first search for quite large constant degree; a more careful calculation may bring down this constant considerably. To conclude the argument for constant expected degree we need the following:

Lemma 3.3 *If l is the biggest number for which $|N_l(s)| \leq n^{2/3}$, the expected value of $\sum_{i=1}^l \frac{1}{|\Gamma_i(s)|}$ is bounded from above by a constant.*

Proof: $|\Gamma_i(s)|$ dominates a random variable with binomial distribution $B(m, p)$, where $m = |\Gamma_{i-1}(s)|(n - n^{-2/3})$. Now, using equation 1.14 in [9], for $\lambda = mp$ and some constants c and c_2 , we have

$$\begin{aligned} E\left(\frac{1}{B(m,p)} \mid B(m,p) > 0\right) &\leq \sum_{k>0} \frac{\exp\{-\lambda - \frac{(\lambda^2+k^2)}{n}\} \lambda^k}{k \cdot k!} \\ &\leq \frac{ce^{-\lambda}}{\lambda} \sum_{k>0} \frac{\lambda^{k+1}}{(k+1)!} \leq \frac{c}{\lambda} \leq \frac{c_2}{|\Gamma_{i-1}(s)|np} \end{aligned}$$

Now

$$\sum_{i=1}^l E\left(\frac{2}{|\Gamma_i(s)|} \mid |\Gamma_i(s)| > 0\right) \leq c_2 \sum_{i=1}^k \frac{1}{(np)^i}$$

\square

We can now proceed as in the proof of Theorem 3.1 to show:

Corollary 3.4 *The expected $v(e, s, t)$ is $O(1)$ in the giant component of G_{np} for $np > 96 \ln 2$.*

Finally, we would very much like to show small constant upper bounds for the expected VCG overcharge in the preferential attachment and the prescribed power-law distributed degree models; for the power-law distributed *expected* degree model, it follows rather easily from the results in [17] that the expected VCG overcharge is $O(\log \log n)$.

Acknowledgments

We wish to thank Miki Ajtai and Elchanan Mossel for useful discussions. The first and last authors were supported by NSF under 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 Tárdoš, 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. and 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-Colell, 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.