Analyzing the small world phenomenon using a hybrid model with local network flow

Reid Andersen * Fan Chung * † Linyuan Lu *

Abstract

Randomly generated graphs with power law degree distribution are typically used to model large real-world networks. These graphs have small average distance. However, the small world phenomenon includes both small average distance and the clustering effect, which is not possessed by random graphs. Here we use a hybrid model which combines a global graph (a random power law graph) with a local graph (a graph with high local connectivity defined by network flow). We present an efficient algorithm which extracts a local graph from a given realistic network. We show that the hybrid model is robust in the sense that for any graph generated by the hybrid model, the extraction algorithm approximately recovers the local graph.

1 Introduction

The small world phenomenon usually refers to two distinct properties — *small average distance* and the *clustering effect*— that are ubiquitous in realistic networks. An experiment by Stanley Milgram [36] titled "The small world problem" indicated that any two strangers are linked by a short chain of acquaintances. The clustering effect implies that any two nodes sharing a neighbor are more likely to be adjacent.

To model networks with the small world phenomenon, one approach is to utilize randomly generated graphs with power law degree distribution. This is based on the observations by several research groups that numerous networks, including Internet graphs, call graphs and social networks, have a *power law* degree distribution, where the fraction of nodes with degree k is proportional to $k^{-\beta}$ for some positive exponent β [1, 2, 5, 6, 7, 11, 13, 18, 23, 27, 29, 31, 37, 40]. Indeed, a random power law graph has small average distances and small diameter. It was shown in [16] that a random power law graph with exponent β, where $2 < \beta < 3$, almost surely has average distance of order $\log\log n$ and has diameter of order $\log n$. (Here, average distance is the average of the distances between pairs of nodes that are connected, and the diameter is the maximum distance between connected pairs.)

In contrast, the clustering effect in realistic networks is often determined by local connectivity and is not amenable to modeling using random graphs. A previous approach to modelling the small world phenomenon was to add random edges to an underlying graph like a grid graph (see Watts and Strogatz [38, 39]). Kleinberg [28] introduced a model where an underlying grid graph G was augmented by random edges placed between each node u, v with probability proportional to $[d_G(u, v)]^{-r}$ for some constant r. In Kleinberg's model and the model of Watts and Strogatz, the subgraphs formed by the random edges have the same expected degree at every node and do not have a power law degree distribution. Fabrikant, Koutsoupias and Paradimitriou [22] proposed a model where vertices are coordinates in the Euclidean plane and edges are added by optimizing

*University of California San Diego, randerse@math.ucsd.edu
†Research supported in part by NSF Grants DMS 0100472 and ITR 0205061

the trade-off between Euclidean distances and "centrality" in the network. Such grid-based models are quite restrictive and far from satisfactory for modeling webgraphs or biological networks.

In [17] Chung and Lu proposed a general hybrid graph model which consists of a global graph (a random power law graph) and a highly connected local graph. The local graph has the property that the endpoints of every edge are joined by at least l edge-disjoint paths each of length at most k, for some fixed parameters k and l. It was shown that these hybrid graphs have average distance and diameter of order $O(\log n)$ where n is the number of vertices.

In this paper, we consider a new notion of local connectivity that is based on network flow. Unlike the problem of finding short disjoint paths, the local flow connectivity can be easily computed using techniques for the general class of fractional packing problems. The goal is to partition a given real-world network into a global subgraph consisting of "long edges" providing small distances and a local graph consisting of "short edges" providing local connections. In this paper, we give an efficient algorithm which extracts a highly connected local graph from any given real world network. We demonstrate that such recovery is robust if the real world graph can be approximated by a random hybrid graph. Namely, we prove that if G is generated by the hybrid graph model, our partition algorithm will recover the original local graph with a small error bound.

2 Preliminaries

2.1 Random graphs with given expected degrees

We consider a class of random graphs with given expected degree sequence $\mathbf{w} = (w_1, w_2, \ldots, w_n)$. The probability that there is an edge between any two vertices v_i and v_j is $p_{ij} = w_i w_j \rho$, where $\rho = (\sum w_i)^{-1}$. We assume that $\max_i w_i^2 < \sum_k w_k$ so that $p_{ij} \leq 1$ for all i and j. It is easy to check that vertex v_i has expected degree w_i. We remark that the assumption $\max_i w_i^2 < \sum_k w_k$ implies that the sequence w_i is graphical [20], except that we do not require the $\{w_i\}$ to be integers. We note that this model allows a non-zero probability for self-loops. The expected number of loops is quite small (of lower order) in comparison with the total number of edges.

We denote a random graph with a given expected degree sequence $\mathbf{w}$ by $G(\mathbf{w})$. For example, the typical random graph $G(n, p)$ (see [21]) on n vertices with edge probability p is just a random graph with expected degree sequence $\mathbf{w} = (pn, pn, \ldots, pn)$.

For a subset S of vertices, we define $\mathrm{Vol}(S) = \sum_{v_i \in S} w_i$ and $\mathrm{Vol}(G) = \sum w_i$. Also for $k \geq 1$, we define $\mathrm{Vol}_k(S) = \sum_{v_i \in S} w_i^k$. We let d denote the average degree $Vol(G)/n$, and let $\tilde{d}$ denote the second order average degree $\mathrm{Vol}_2(G)/\mathrm{Vol}(G)$.

2.2 Random power law graphs

A random power law graph $M(n, \beta, d, m)$ is a random graph $G(\mathbf{w})$ whose expected degree sequence $\mathbf{w}$ is determined by the following four parameters.

- n is the number of vertices.
- $\beta > 2$ is the power law exponent.
- d is the expected average degree.

- m is the maximum expected degree and $m^2 = o(nd)$.

We remark that most realistic graphs have degree sequence satisfying the power law for a certain range of degrees (not too small or too large). An alternative definition for m is the maximum within the range of degrees that the power law holds.

We let the i-th vertex v_i have expected degree

$$w_i = ci^{-\frac{1}{\beta-1}}$$

for $i \geq i_0$, some c and i_0 (to be chosen later). It is easy to compute that the number of vertices of expected degree between k and $k+1$ is of order $c'k^{-\beta}$ where $c' = c^{\beta-1}(\beta-1)$, as required by the power law. To determine c, we consider

$$\begin{aligned} \mathrm{Vol}(G) &= \sum_i w_i = \sum_{i=i_0}^{n} ci^{\frac{1}{\beta-1}} \\ &\approx c\frac{\beta-1}{\beta-2}n^{1-\frac{1}{\beta-1}} \end{aligned}$$

Here we assume $\beta > 2$. Since $nd \approx \mathrm{Vol}(G)$, we choose

$$c = \frac{\beta-2}{\beta-1}dn^{\frac{1}{\beta-1}} \tag{1}$$

$$i_0 = n(\frac{d(\beta-2)}{m(\beta-1)})^{\beta-1} \tag{2}$$

Let $f(x) = cx^{-\frac{1}{\beta-1}}$. The expected degrees (or weights) are just $f(i)$, $i_0 \leq i \leq n$.

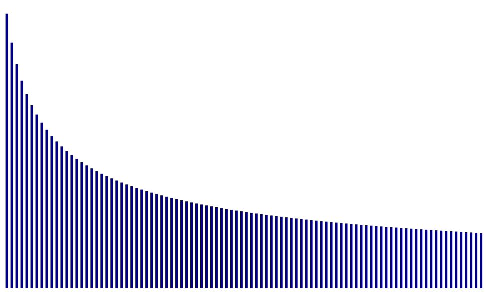

Figure 1: *Weight distribution $f(x)$.*

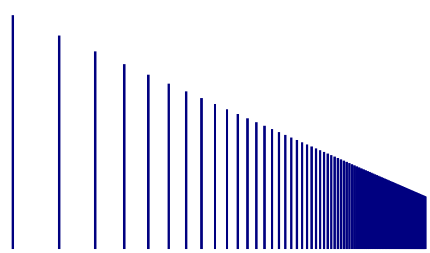

Figure 2: *Log-scale of figure 1.*

2.3 A concentration inequality

Let $X_1, \ldots, X_n$ be independent random variables with

$$Pr(X_i = 1) = p_i, \qquad Pr(X_i = 0) = 1 - p_i$$

For $X = \sum_{i=1}^n a_i X_i$, we have $E(X) = \sum_{i=1}^n a_i p_i$ and we define $\nu = \sum_{i=1}^n a_i^2 p_i$. Then we have (see [15]).

$$Pr(X < E(X) - \lambda) \leq e^{-\lambda^2/2\nu} \tag{3}$$

$$Pr(X > E(X) + \lambda) \leq e^{-\frac{\lambda^2}{2(\nu + a\lambda/3)}} \tag{4}$$

where $a = \max\{a_1, a_2, \ldots, a_n\}$.

2.4 A bound for sums

We will use the following trivial but convenient bound.

Lemma 1 *Let X be some finite set with nonnegative weights $w(x)$, and let $A \subseteq X^k$ be a set of ordered k-tuples from X. If each element $x \in X$ appears in at most M elements of A, then*

$$\sum_{(x_{i_1} \dots x_{i_k}) \in A} w(x_{i_1}) \cdots w(x_{i_k}) \leq M \sum_{x \in X} w(x)^k$$

Proof: Order the elements of X as $x_1 \dots x_n$ such that $w(x_1) \geq \cdots \geq w(x_n)$. Let A_j be the collection of tuples $v \in A$ where j is the smallest index of any element in v. We have $|A_j| \leq M$, and $\cup_{j \in [1,n]} A_j = A$, so

$$\begin{aligned} \sum_{(x_{i_1} \dots x_{i_k}) \in A} w(x_{i_1}) \cdots w(x_{i_k}) &\leq \sum_{j \in [1,n]} \sum_{(x_{i_1} \dots x_{i_k}) \in A_j} w(x_{i_1}) \cdots w(x_{i_k}) \\ &\leq \sum_{j \in [1,n]} M w(x_j)^k \\ &\leq M \sum_{x \in X} w(x)^k \end{aligned}$$

3 Local graphs and hybrid graphs

3.1 Local graphs

There are a number of ways to define local connectivity between two given vertices u and v. A natural approach is to consider the maximum number $a(u,v)$ of short edge-disjoint paths between the vertices, where *short* means having length at most ℓ. Another approach is to consider the minimum size $c(u,v)$ of a set of edges whose removal leaves no short path between the vertices. When we restrict to short paths, the analogous version of the max-flow min-cut theorem does not hold, and in fact a and c can be different by a factor of $O(\ell)$ ([44],[45]). However we still have the trivial relations $a \leq c \leq \ell \cdot a$.

Both of the above notions of local connectivity are difficult to compute, and in fact computing the maximum number of short disjoint paths is $\mathcal{NP}$-hard if $\ell \geq 4$ [43]. Instead we will consider the maximum short flow between u and v. The maximum short flow can be computed in polynomial time using nontrivial but relatively efficient algorithms for fractional packing (see Section 5). The most compelling reason for using short flow as a measure of local connectivity is that it captures the spirit of $a(u,v)$ and $c(u,v)$, but is efficiently computable.

Formally, a short flow is a positive linear combination of short paths where no edge carries more than 1 unit of flow. Finding the maximum short flow can be viewed as a linear program. Let P_ℓ be the collection of short u-v paths, and let P_e be the collection of short u-v paths which intersect the edge e.

Definition 1 (Flow connectivity) *A short flow is a feasible solution to the following linear program. The flow connectivity $f(u,v)$ between two vertices is the maximum value of any short flow, which is the optimum value of the following LP problem:*

$$
\begin{aligned}
\text{maximize} \quad & \sum_{p \in P_\ell} f_p & & (5)\\
\text{subject to} \quad & \sum_{p \in P_e} f_p \le 1 & \text{for each } e \in L &\\
& f_p \ge 0 & \text{for each } p \in P_\ell &
\end{aligned}
$$

The linear programming dual of the flow connectivity problem is a fractional cut problem: to find the minimum weight cutset so that every short path has at least 1 unit of cut weight. This gives us the following relation between a, c, and f.

$$a \le f \le c \le a \cdot \ell$$

We say two vertices u and v are (f, ℓ)-connected if there exists an short flow between them of size at least f.

Definition 2 (Local Graphs) *A graph L is an (f, ℓ)-local graph if for each edge $e = (u, v)$ in L, the vertices u and v are (f, ℓ)-connected in $L \setminus \{e\}$.*

3.2 Hybrid power law graphs

A hybrid graph H is the union of the edge sets of an (f, ℓ)-local graph L and a global graph G on the same vertex set. We here consider the case where the global graph $G(\mathbf{w})$ is a power law graph $M(n, \beta, d, m)$. The following theorems were proved in [17].

Theorem 1 *For a hybrid graph H with $G = M(n, \beta, d, m, L)$ and $\beta > 3$, almost surely, the average distance is $(1 + o(1))\frac{\log n}{\log \tilde{d}}$ and the diameter is $O(\log n)$.*

Theorem 2 *For a hybrid graph H with $G = M(n, \beta, d, m, L)$ and $2 < \beta < 3$, almost surely, the average distance is $O(\log \log n)$ and the diameter is $O(\log n)$. For a hybrid graph H with $G = M(n, \beta, d, m, L)$ and $\beta = 3$, almost surely, the average distance is $O(\log n / \log \log n)$ and the diameter is $O(\log n)$.*

For the range of $2 < \beta < 3$, the power law graphs include many real networks. We can further reduce the diameter if additional conditions are satisfied. A local graph L is said to have isoperimetric dimension δ if for every vertex v in L and every integer $k < (\log \log n)^{1/\delta}$, there are at least k^δ vertices in L of distance k from v. For example, the grid graph in the plane has isoperimetric dimension 2. The d-dimensional grid graph has isoperimetric dimension d.

Theorem 3 *In a hybrid graph H with $G = M(n, \beta, d, m, L)$ and $2 < \beta < 3$, suppose that the local graph has isoperimetric dimension δ, where $\delta \ge \log \log n / (\log \log \log n)$. Then almost surely, the diameter is $O(\log \log n)$.*

Theorem 4 *In a hybrid graph H with $G = M(n, \beta, d, m, L)$ and $2 < \beta < 3$, suppose that the local graph has isoperimetric dimension δ. Then almost surely, the diameter is $O((\log n)^{1/\delta})$.*

Theorem 5 *In a hybrid graph H with $G = M(n, \beta, d, m, L)$ and $2 < \beta < 3$, suppose that every vertex is within distance $\log \log n$ of some vertex of degree $\log n$. Then almost surely, the diameter is $O(\log \log n)$.*

4 Extracting the local graph

For a given graph, the problem of interest is to extract the largest (f,ℓ)-local subgraph. We define $L_{f,\ell}(G)$ to be the union of all (f,ℓ)-local subgraphs in H. By definition, the union of two (f,ℓ)-local graphs is an (f,ℓ)-local graph, and so $L_{f,\ell}(G)$ is in fact the unique largest (f,ℓ)-local subgraph in G. We remark that $L_{f,\ell}(G)$ is not necessarily connected. There is a simple greedy algorithm to compute $L_{f,\ell}(G)$ in any graph G.

4.1 An algorithm for extracting the local graph

Extract(f,ℓ)**:** We are given as input a graph G and parameters (f,ℓ). Let $H = G$. If there is some edge $e = (u,v)$ in H where u and v are not (f,ℓ)-connected in $H \setminus \{e\}$, then let $H = H \setminus \{e\}$. Repeat until no further edges can be removed, then output H.

Theorem 6 *For any graph G and any (f,ℓ),* **Extract**(f,ℓ) *returns $L_{f,\ell}(G)$.*

Proof: Given a graph G, let L' be the graph output by the greedy algorithm. A simple induction argument shows that each edge removed by the algorithm is not part of any (f,l)-local subgraph of G, and thus $L_{f,\ell}(G) \subseteq L'$. Since no further edges can be removed from L', L' is (f,l)-local and so $L' \subseteq L_{f,\ell}(G)$. Thus $L' = L_{f,\ell}(G)$.

Time analysis: The algorithm requires $O(|E|^2)$ maximum short flow computations. In section 5 we will describe an algorithm to compute the maximum short flow.

4.2 Recovering the local graph

When applied to a hybrid graph $H = G \cup L$ with an (f,ℓ)-local graph L, the algorithm **Extract**(f,ℓ) will output $L_{f,\ell}$, which is almost exactly L if G is sufficiently sparse. Note that $L \subseteq L_{f,\ell}$ by definition of the local graph.

Theorem 7 *Let $H = G \cup L$ be a hybrid graph where L is (f,ℓ)-local with maximum degree M, and where $G = G(w)$ with average weight d, second order average weight $\tilde{d}$, and maximum weight m. Let $L' = L_{f,\ell}(H)$. If $\tilde{d}$ satisfies*

$$\tilde{d} \leq n^{\alpha} \leq \left(\frac{nd}{m^2}\right)^{1/\ell} n^{-3/f\ell} \text{ for some constant } \alpha > 0,$$

Then with probability $1 - O(n^{-1})$:

1. *$L' \setminus L$ contains $O(\tilde{d})$ edges.*
2. *$d_{L'}(x,y) \geq \frac{1}{\ell} d_L(x,y)$ for every pair $x, y \in L$.*

In the special case where all the weights are equal and $G(w) \sim G(n,p)$, Theorem 7 has a cleaner statement, and is tight in the sense that if d is larger than $n^{\frac{1}{\ell}}$ we cannot hope to recover a good approximation to the original local graph.

Theorem 8 *Let H be a hybrid graph as in Theorem 7 and let $G = G(n,p)$ with $p = dn^{-1}$. If*

$$d \leq n^{\alpha} \leq n^{1/\ell} n^{-3/f\ell} \text{ for some constant } \alpha > 0,$$

Then with probability $1 - O(n^{-1})$, results (1)-(2) from Theorem 7 hold.

Theorem 9 *Let G be chosen from $G(n,p)$ where $p = dn^{-1}$, and let*

$$d \geq 6fn^{\frac{1}{\ell}}(\log n)^{\frac{1}{\ell}}.$$

With probability $1 - O(n^{-2})$, $L_{f,\ell}(H) = H$.

We also point out that the term $\left(\frac{nd}{m^2}\right)^{1/\ell}$ in Theorem 7 is nearly optimal, although we will not make this precise. In the $G(w)$ model, $\tilde{d}$ replaces d since it is roughly the factor we expect a small neighborhood to expand at each step, and having some dependence on m is unavoidable.

5 Computing the maximum short flow

The problem of finding the maximum short flow between u and v can be viewed as a fractional packing problem, as introduced by Plotkin, Shmoys, and Tardos [41]. A fractional packing problem has the form

$$\max\{\, \mathbf{c}^{\mathbf{T}}\mathbf{x} \mid A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \succeq \vec{\mathbf{0}}\,\}.$$

To view the maximum short flow as a fractional packing problem, first let $G(u,v)$ be a subgraph containing all short paths from u to v. For example, we may take $G(u,v) = N_{\ell/2}(u) \cup N_{\ell/2}(v)$. Let A be the incidence matrix where each row represents an edge in $G(u,v)$ and each column represents a short path from u to v. Let $\mathbf{b} = \mathbf{c} = \vec{\mathbf{1}}$.

The algorithm **Max Short Flow** below is an implementation of the fractional packing algorithm by Garg and Könemann [42] which has been specialized for our problem. See their paper for more details and a proof of correctness.

Algorithm MAX SHORT FLOW:
Compute $E(u,v) = N_{\ell/2}(u) \cup N_{\ell/2}(v)$ using breadth-first search.
Set $w(e) = 1$ for each $e \in G(u,v)$. %w(e) is an edge-weight
Repeat until $\sum_e w(e) \geq (1+\epsilon)^{-1}\left((1+\epsilon)m\right)^{1/\epsilon}$:
  Let p be the short path minimizing $\sum_{e \in p} w(e)$. %Compute using **MINWEIGHT**.
  Set $\alpha(p) = \alpha(p) + 1$ % Route 1 additional unit on p. Initially $\alpha(p) = 0$ implicitly.
  Set $f = f + 1$ % Record that we have augmented the flow by 1.
  For each $e \in p$:
    Set $w(e) = w(e)(1+\epsilon)$. %Increase weights on edges in p
    Set $c(e) = c(e) + 1$. %Record increased congestion on edges in p
Then to conclude:
  Let $C = \max_e c(e)$ %C is the the maximum congestion.
  Set $\alpha(p) = (1/C)\alpha(p)$ for all p % Scale to obtain a feasible flow.
  Output $\{\alpha(p)\}$ and f/C %(Output the flow and its value)

Algorithm MINWEIGHT:
Let $S_0 = \{u\}$ and let $\phi_0(u) = 0$.
Repeat for $k \in [1, \ell]$:
  Let $S_k = S_{k-1} \cup \Gamma(S_{k-1})$ in $G(u, v)$
  For each $x \in S_k$:
    Let $\phi_k(x) = \min_{y \in \Gamma(x)}(w(xy) + \phi_{k-1}(y))$.
    Let $\psi_k(v)$ be some vertex which minimizes this quantity for v at step k.
% We will now reconstruct the minimum-weight path
Find the index k of the minimum value among $\phi_0(v) \ldots \phi_\ell(v)$.
Let $v = v_k$.
For $j = k \ldots 1$
  Let $v_{j-1} = \psi_j(v_j)$
Output $v_0 \ldots v_k$

Minweight runs in time $T_{Min} = O(m\ell)$ where m is the number of edges in $G(u, v)$.

Theorem 10 **MAX SHORT FLOW** *produces a* $(1-\epsilon)^{-2}$*-approximation to the maximum short flow in time* $O(m\lceil \frac{1}{\epsilon} log_{1+\epsilon} m \rceil) T_{Min} = O(m^2 \ell \lceil \frac{1}{\epsilon} log_{1+\epsilon} m \rceil)$*, where* m *is the number of edges in* $G(u, v)$.

This follows from the work of Garg and Könemann in [42].

Experiment: We have implemented the **Extract** algorithm and tested it on various graphs. Smaller values of ϵ give a more accurate output, but with a longer running time. For some hybrid graphs, the local graphs are almost perfectly recovered (see figure 4).

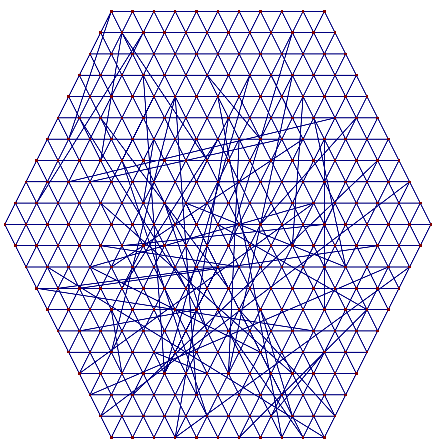

Figure 3: *A hybrid graph, containing the hexagonal grid graph as a local graph.*

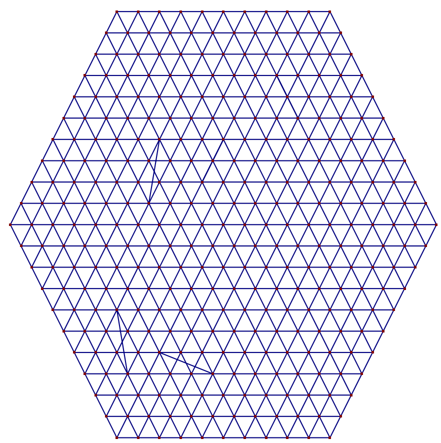

Figure 4: *After applying* **Extract** *(with parameters* $k = 2.5$, $l = 4$ *and* $\epsilon = 0.5$*), the local graph is almost perfectly recovered.*

6 Communities and examples

In this paper, we have used the hybrid graph model to understand the "landscape" of many real world graphs. We showed that the **Extract** algorithm is able to extract the local structures, which are rarely found in the random graphs. The local graph L found by the **Extract**(f, l) algorithm may not be connected. Each connected component L can be viewed as a local community. In another words, we define a "community" of a vertex v to be the maximum subgraph containing v, which is connected and locally (f, l) connected. By using different parameters f and l, we have a hierarchy of communities.

There is a large literature on clustering data into communities [19, 25]. For example, Flake et al. [26] define communities using minimum cut trees. Communities in our definition are monotone properties. I.e., adding an edge to the graph G increases the sizes of communities. The communities found by our **Extract** algorithm often have rich structures other than cliques or complete bipartite subgraphs.

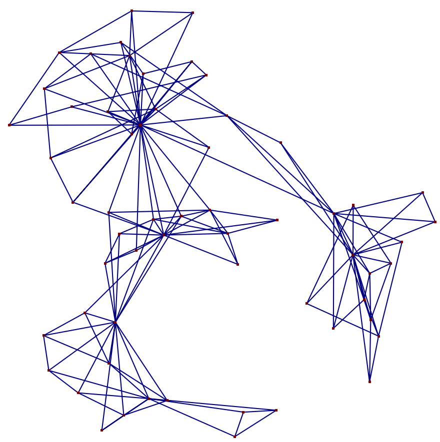

Figure 5: *A community of size 25 in a routing graph, which is* $(3, 3)$ *local connected.*

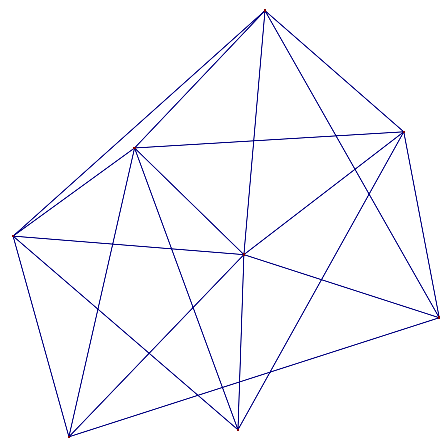

Figure 6: *A sub-community, which is* $(4, 3)$ *locally connected, sits inside the community of size 35.*

An example: We used the **Extract** algorithm on a routing graph G using data collected by "champagne.sdsc.org", having 9175 vertices and 15519 edges. The maximum 3-connected subgraph of G consists of 7 K_4's and a large connected component L with 2364 vertices and 5947 edges. Our **Extract**$(3, 3)$ algorithm breaks L into 79 non-trivial communities of various sizes. The largest community has 881 vertices. The second largest community (of size 59) is illustrated in Figure 5 and two communities of size 25 and 35 are illustrated in Figure 7 and 8.

References

[1] L. A. Adamic and B. A. Huberman, Growth dynamics of the World Wide Web, *Nature*, **401**, September 9, 1999, pp. 131.

[2] W. Aiello, F. Chung and L. Lu, A random graph model for massive graphs, *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, (2000) 171-180.

[3] W. Aiello, F. Chung and L. Lu, Random evolution in massive graphs, Extended abstract appeared in *The 42th Annual Symposium on Foundation of Computer Sciences*, October, 2001. Paper version appeared in *Handbook on Massive Data Sets*, (Eds. J. Abello, et. al.), Kluwer Academic Publishers (2002), 97-122.

[4] N. Alon and J. H. Spencer, *The Probabilistic Method*, Wiley and Sons, New York, 1992.

[5] R. B. R. Azevedo and A. M. Leroi, A power law for cells, *Proc. Natl. Acad. Sci. USA*, vol. **98**, no. 10, (2001), 5699-5704.

[6] Albert-László Barabási and Réka Albert, Emergence of scaling in random networks, Science **286** (1999) 509-512.

[7] A. Barabási, R. Albert, and H. Jeong, Scale-free characteristics of random networks: the topology of the world wide web, *Physica* A 272 (1999), 173-187.

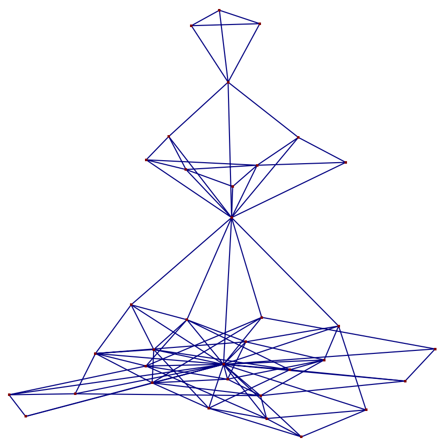

Figure 7: *A community of size 35 in the routing graph.*

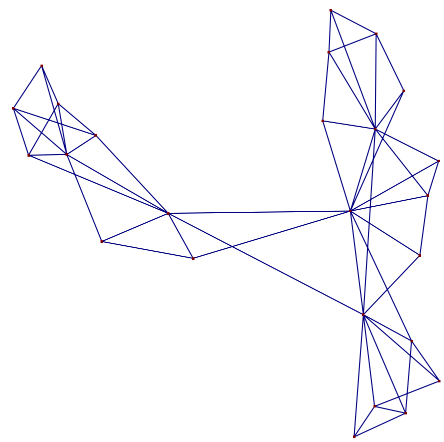

Figure 8: *A community of size 25 in the routing graph.*

[8] E. A. Bender and E. R. Canfield, The asymptotic number of labelled graphs with given degree sequences, *J. Combinat. Theory (A)*, **24**, (1978), 296-307.

[9] B. Bollobás, *Random Graphs*, Academic, New York, 1985.

[10] B. Bollobás, O. Riordan, J. Spencer and G. Tusnády, The Degree Sequence of a Scale-Free Random Graph Process, Random Structures and Algorithms, Vol. **18**, no. 3 (2001), 279-290.

[11] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tompkins, and J. Wiener, "Graph Structure in the Web," *proceedings of the WWW9 Conference*, May, 2000, Amsterdam. Paper version appeared in *Computer Networks* **33**, (1-6), (2000), 309-321.

[12] CAIDA, the Cooperative Association for Internet Data Analysis, http://www.caida.org/

[13] K. Calvert, M. Doar, and E. Zegura, Modeling Internet topology. *IEEE Communications Magazine*, **35(6)** (1997) 160-163.

[14] Fan Chung and Linyuan Lu, The diameter of random sparse graphs, *Advances in Applied Math.*, **26** (2001), 257-279.

[15] Fan Chung and Linyuan Lu, Connected components in a random graph with given degree sequences, *Annals of Combinatorics*, **6** (2002), 125-145.

[16] F. Chung and L. Lu, Average distances in random graphs with given expected degree sequences, *Proceedings of National Academy of Science*, **99** (2002), 15879-15882.

[17] F. Chung and L. Lu, The small world phenomenon in hybrid power law graphs *Lecture Note in Physics* speical volume on "Complex Network", to appear.

[18] C. Cooper and A. Frieze, On a general model of web graphs, *Random Structures and Algorithms* Vol. **22**, (2003), 311-335.

[19] B. Everitt, **Clustering analysis.** Halsted Press, New York, 1980.

[20] P. Erdős and T. Gallai, Gráfok előírt fokú pontokkal (Graphs with points of prescribed degrees, in Hungarian), *Mat. Lapok* **11** (1961), 264-274.

[21] P. Erdős and A. Rényi, On random graphs. I, *Publ. Math. Debrecen* **6** (1959), 290-291.

[22] A. Fabrikant, E. Koutsoupias and C. H. Papadimitriou, Heuristically optimized trade-offs: a new paradigm for power laws in the Internet, *STOC* 2002.

[23] M. Faloutsos, P. Faloutsos, and C. Faloutsos, On power-law relationships of the Internet topology, *Proceedings of the ACM SIGCOM Conference*, Cambridge, MA, 1999.

[24] M. Henzinger, private communication.

[25] G. W. Flake, S. Lawrence, and C. L. Giles. Efficient identification of web communities. In *Proceedings of the Sixth International Conference on Knowledge Discovery and Data Mining*, Boston, MA, 2000.

[26] G. W. Flake, R. E. Tarjan, and K. Tsioutsiouliklis, Graph Clustering and Minimum Cut Trees.

[27] S. Jain and S. Krishna, A model for the emergence of cooperation, interdependence, and structure in evolving networks, *Proc. Natl. Acad. Sci. USA*, vol. **98**, no. 2, (2001), 543-547.

[28] J. Kleinberg, The small-world phenomenon: An algorithmic perspective, *Proc. 32nd ACM Symposium on Theory of Computing*, 2000.

[29] J. Kleinberg, S. R. Kumar, P. Raghavan, S. Rajagopalan and A. Tomkins, The web as a graph: Measurements, models and methods, *Proceedings of the International Conference on Combinatorics and Computing*, 1999.

[30] S. R. Kumar, P. Raghavan, S. Rajagopalan and A. Tomkins, Trawling the web for emerging cyber communities, *Proceedings of the 8th World Wide Web Conference*, Toronto, 1999.

[31] S. R. Kumar, P. Raghavan, S. Rajagopalan and A. Tomkins, Extracting large-scale knowledge bases from the web, *Proceedings of the 25th VLDB Conference*, Edinburgh, Scotland, 1999.

[32] Linyuan Lu, The diameter of random massive graphs, *Proceedings of the Twelfth ACM-SIAM Symposium on Discrete Algorithms*, (2001) 912-921.

[33] C. McDiarmid, Concentration. Probabilistic methods for algorithmic discrete mathematics, *Algorithms Combin.*, **16**, Springer, Berlin, (1998) 195–248.

[34] Michael Molloy and Bruce Reed, A critical point for random graphs with a given degree sequence. *Random Structures and Algorithms*, Vol. **6**, no. 2 and 3 (1995), 161-179.

[35] Michael Molloy and Bruce Reed, The size of the giant component of a random graph with a given degree sequence, *Combin. Probab. Comput.* **7**, no. 3 (1998), 295-305.

[36] S. Milgram, The small world problem, *Psychology Today*, **2** (1967), 60-67.

[37] M. E. J., Newman, The structure of scientific collaboration networks, *Proc. Natl. Acad. Sci. USA*, vol. **98**, no. 2, (2001), 404-409.

[38] D. J. Watts, *Small Worlds — The Dynamics of Networks between Order and Randomness*, Princeton University Press, New Jersey, 1999.

[39] D. J. Watts and S. H. Strogats, Collective dynamics of 'small world' networks, *Nature* **393**, 440-442.

[40] E. Zegura, K. Calvert, and M. Donahoo, A quantitative comparison of graph-based models for Internet topology. *IEEE/ACM Transactions on Networking*, **5** (6), (1997), 770-783.

[41] S. Plotkin, D. B. Shmoys, and E Tardos, Fast approximation algorithms for fractional packing and covering problems, *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science*, 1991, pp. 495–504. http://citeseer.ist.psu.edu/plotkin95fast.html

[42] N. Garg, J. Konemann, Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *Technical Report, Max-Planck-Institut fur Informatik, Saarbrucken, Germany* (1997). http://citeseer.ist.psu.edu/garg97faster.html

[43] A. Itai, Y. Perl, and Y. Shiloach, The complexity of finding maximum disjoint paths with length constraints, *Networks* **12** (1982)

[44] L. Lovász, V. Neumann-Lara, M. Plummer, Mengerian theorems for paths of bounded length, *Periodica Mathematica Hungaria* **9** (1978)

[45] S. Boyles, G. Exoo, On line disjoint paths of bounded length. *Discrete Math.* **44** (1983)

7 Appendix

7.1 Proof of Theorem 7

We say an edge in H is *global* if it is in $G \setminus L$. A global edge is *long* if $d_L(u,v) > \ell$ and *short* otherwise. We will show that under the hypotheses of Theorem 7 no long edges are likely to survive. If (u,v) is a short edge in G, it is possible that there is a short flow of size f from u to v entirely through edges in L. This means we can not say a short edge is unlikely to survive without placing additional assumptions on the local graph. However, an easy computation shows there are not likely to be many short edges in G.

Lemma 2 *The expected number of short edges in $L_{f,\ell}(G) \setminus L$ is $O(\tilde{d})$.*

The expected number of short edges in $L_{f,\ell}(G) \setminus L$ is

$$\begin{aligned}
\sum_{\text{short } (x,y)} \Pr\left[(x,y) \in L_{f,\ell}(G)\right] &\leq \sum_{\text{short } (x,y)} \Pr[(x,y) \in G] \\
&\leq \sum_{\substack{(x,y) \\ x \in G, y \in N_\ell^L(x)}} w_x w_y \rho \\
&\leq 2M^\ell \sum_{x \in G} w_x^2 \rho \qquad (*) \\
&= 2M^\ell \tilde{d} \\
&= O(\tilde{d})
\end{aligned}$$

The line marked (*) is obtained by applying Lemma 1 to the sum in the previous line, noting that each vertex x appears in at most $2M^\ell$ terms.

Proposition 1 *The probability that a given long edge survives is $O(n^{-3})$.*

This is the more difficult part of the proof and will take some work.

Definition 3 *$N_k(u)$ and $\Gamma_k(u)$*

For $k \in [0,\ell]$, let $N_k(u)$ be the set of vertices y such that there exists a path $p = v_0 \dots v_k$ from u to y in H obeying the following condition:

$$d_L(v_i, v) > \ell - i \text{ for all } i \in [0,k].$$

We define $\Gamma_k(u)$ to be the corresponding strict neighborhood,

$$\Gamma_k(u) = \{\, y \mid y \in N_k(u), y \notin N_0(u) \dots N_{k-1}(u) \,\}.$$

The following recursive definition of $\Gamma_k(u)$ will be useful, and is easily seen to be equivalent to the original.

$$\begin{aligned}
\Gamma_0(u) &= \{u\} \\
\Gamma_k(u) &= \left\{ y \;\middle|\; \begin{array}{l} y \notin N_{\ell-k}(v), \\ y \notin \Gamma_0(u) \dots \Gamma_{k-1}(u), \\ (x,y) \in H \text{ for some } x \in \Gamma_{k-1}(u) \end{array} \right\}
\end{aligned}$$

Definition 4 $C(u,v)$

Define $C(u,v)$ to be the set of edges in

$$\bigcup_{k\in[1,\ell]} \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right).$$

Remark 1 *All the edges in C(u,v) are global edges.*

If $(x,y) \in \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right)$, then $d_L(x,v) > \ell - (k-1)$ and $d_L(y,v) \leq \ell - k$. Thus $d_L(x,y) \geq 2$, so (x,y) cannot be a local edge and must be global.

Lemma 3 *Let (u,v) be a surviving long edge in H. Then $f \leq |C(u,v)|$.*

Proof of Lemma 3: We first show that every short path between u and v in H contains an edge from $C(u,v)$. Let $p = v_0 \dots v_k$ be a path of length $k \leq \ell$ between u and v in H. The last vertex on the path is $v_k = v$, so we have $d_L(v_k, v) = 0 \leq \ell - k$, and thus $v_k \notin N_k(u)$. The first vertex on the path is $v_0 = u$, and thus $v_0 \in N_0(u)$. Let $j \geq 1$ be the smallest integer such that $v_j \notin N_j(u)$. By definition, $v_{j-1} \in N_{j-1}(u)$, while $v_j \notin N_j(u)$. This implies $d_L(v_j, v) \leq \ell - j$, so $v_j \in N^L_{\ell-j}(v)$. We now have that $v_{j-1}v_j$ is an edge in $N_{j-1}(u) \times N^L_{\ell-j}(v)$. We conclude that $v_{j-1}v_j$ is an edge in $C(u,v)$ by noticing

$$N_{j-1}(u) \times N^L_{\ell-j}(v) \subseteq \bigcup_{k\in[1,j]} \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right).$$

We can now complete the proof. If the set $C(u,v)$ is removed, then no short paths remain between u and v. Thus, if a is the maximum number of short disjoint paths, c is the size of the minimum cut to remove all short paths, and f is the maximum ℓ-flow, we have

$$a \leq f \leq c \leq |C(u,v)|.$$

Thus $f \leq |C(u,v)|$, and in fact the lemma would also hold if we were considering disjoint paths or cuts.

Lemma 4 *If we condition on the values of the sets $\Gamma_0(u) \dots \Gamma_{\ell-1}(u)$, then the edges in*

$$\bigcup_{k\in[1,\ell]} \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right).$$

are mutually independent and occur with the same probabilities as in G.

Proof of Lemma 4: We will reveal $\Gamma_0(u) \dots \Gamma_{\ell-1}(u)$ sequentially by a breadth-first search. From the recursive definition of $\Gamma_k(u)$, it is clear that we can determine $\Gamma_k(u)$ given $\Gamma_{k-1}(u)$ by examining only the edges in

$$\Gamma_{k-1} \times \left(V \setminus (N^L_{\ell-k} \cup \Gamma_0(u) \cup \dots \cup \Gamma_{k-1}(u))\right).$$

In particular, in determining $\Gamma_k(u)$ from $\Gamma_{k-1}(u)$ we do not examine any edges with an endpoint in $\Gamma_0(u) \dots \Gamma_{k-2}(u)$, and we do not examine any edges in $\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)$. Thus, we do not examine any edges in

$$\bigcup_{j\in[1,k]} \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right)$$

when revealing $\Gamma_k(u)$, and we may reveal $\Gamma_0(u)\dots\Gamma_{\ell-1}(u)$ without examining any edges in

$$\bigcup_{k\in[1,\ell]} \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right).$$

Lemma 5 *With probability* $1 - e^{-\Omega(n^\alpha)}$,

$$\sum_{k\in[1,\ell]} Vol(\Gamma_{k-1}(u))Vol(N^L_{\ell-k}(v)) \leq 4m^2(4M\hat{d})^{\ell-1}$$

Let G_j denote the set of global edges between $\Gamma_{j-1}(u)$ and $\Gamma_j(u)$. Let $\Gamma_{k,j}(u)$ be the set of vertices $x \in \Gamma_k(u)$ where there exists a path $p = v_0 \dots v_k$ from u to x where $v_i \in \Gamma_i(u)$ and j is the minimum number such that $v_j \dots v_k$ consists entirely of local edges. Thus, if $j \neq 0$ the edge $v_{j-1}v_j$ is in G_j, and if $j = 0$ then we have $x \in \Gamma_{k,0}(u) = \Gamma^L_k(u)$. We think of $\Gamma_{k,j}(u)$ as the collection of vertices in $\Gamma_k(u)$ which were guaranteed to be in $\Gamma_k(u)$ by an edge in G_j, or in the case $j = 0$ were guaranteed to be in $\Gamma_k(u)$ by being in $N^L_k(u)$. We will be considering the volumes of these sets, so we define

$$\begin{aligned} V_k &= Vol(\Gamma_k(u)) \\ V_{k,j} &= Vol(\Gamma_{k,j}(u)), \end{aligned}$$

and make note of the following simple facts:

$$\begin{aligned} \Gamma_k(u) &= \bigcup_{j\in[0,k]} \Gamma_{j,k}(u) \\ V_k &\leq \sum_{j\in[0,k]} V_{k,j}. \end{aligned}$$

We will now give an upper bound on $V_{k,j}$ conditional on V_{j-1}.

Proposition 2 *Let* $\hat{V}_j = \max\{V_j, m\}$*, and* $\hat{d} = n^\alpha \geq \tilde{d}$*. With probability* $1 - exp(-\Omega(n^\alpha))$,

$$V_{k,j} \leq \left(4M^{k-j}\hat{d}\right)\hat{V}_{j-1} \qquad \textit{for all } j \leq k \leq \ell - 1$$

Proof of Proposition 2:

$$\begin{aligned} V_{k,j} &= Vol\left(\bigcup_{(x,y)\in G_j} \Gamma^L_{k-j}(y)\right) \\ &\leq \sum_{(x,y)\in G_j} Vol(\Gamma^L_{k-j}(y)) \end{aligned}$$

We wish to bound this quantity, so we define

$$Y_{k,j} = \sum_{(x,y)\in G_j} Vol\left(\Gamma^L_{k-j}(y)\right)$$

We will use the concentration inequality (4) to bound $Y_{k,j}$, so we first compute a, μ, and ν.

$$a(Y_{k,j}) = \max_x \left\{ Vol\left(\Gamma^L_{k-j}(x)\right)\right\} \leq M^{k-j} m$$

$$\begin{aligned}
\mu(Y_{k,j}) &= \sum_{(x,y)} \Pr\left[(x,y) \in G_j\right] Vol\left(\Gamma^L_{k-j}(y)\right) \\
&\leq \sum_{y \in G} E\left[\#\text{Global edges between } \Gamma_{j-1}(u) \text{ and } y\right] \cdot Vol\left(\Gamma^L_{k-j}(y)\right) \\
&= \sum_{y \in G} \left(\sum_{x \in \Gamma_{j-1}(u)} w_x w_y \rho \right) Vol\left(\Gamma^L_{k-j}(y)\right) \\
&= \rho \sum_{y \in G} w_y \left(\sum_{x \in \Gamma_{j-1}(u)} w_x \right) Vol\left(\Gamma^L_{k-j}(y)\right) \\
&= \rho V_{j-1} \sum_{y \in G} w_y Vol\left(\Gamma^L_{k-j}(y)\right) \\
&= \rho V_{j-1} \sum_{\substack{(x,y) \\ y \in G, x \in \Gamma^L_{k-j}(y))}} w_x w_y \\
&\leq \rho V_{j-1} 2M^{k-j} \sum_{y \in G} w_y^2 \qquad (*) \\
&= \left(2M^{k-j}\tilde{d}\right) V_{j-1}
\end{aligned}$$

The line marked (*) is obtained by applying Lemma 1 to the sum in the previous line, noting that each vertex appears in at most $2M^{k-j}$ terms.

$$
\begin{aligned}
\nu(Y_{k,j}) &= \sum_{(x,y)} Vol\left(\Gamma^L_{k-j}(y)\right)^2 \Pr\left[(x,y) \in G_j\right] \\
&\leq \sum_{y \in G} Vol\left(\Gamma^L_{k-j}(y)\right)^2 E\left[\#\text{Global edges between } \Gamma_{j-1}(u) \text{ and } y\right] \\
&\leq \sum_{y \in G} \left(\sum_{(x,z) \in \Gamma^L_{k-j}(y)} w_x w_z \right) \left(\sum_{v \in \Gamma_{j-1}(u)} w_v w_y \rho \right) \\
&\leq \rho V_{j-1} \sum_{y \in G} \left(\sum_{(x,z) \in \Gamma^L_{k-j}(y)} w_x w_z \right) w_y \\
&= \rho V_{j-1} \left(\sum_{\substack{(x,y,z): \\ x \in G, (y,z) \in \Gamma^L_{k-j}(x)}} w_x w_y w_z \right) \\
&\leq \rho V_{j-1} \left(3(M^{k-j})^2\right) \sum_{y \in G} w_y^3 \qquad (*) \\
&\leq V_{j-1} \left(3(M^{k-j})^2\right) m \sum_{y \in G} w_y^2 \rho \\
&= \left(3(M^{k-j})^2 m \tilde{d}\right) V_{j-1}
\end{aligned}
$$

In the line marked (*), we are asserting

$$
\left(\sum_{\substack{(x,y,z): \\ x \in G, (y,z) \in \Gamma^L_{k-j}(x)}} w_x w_y w_z \right) \leq \left(3(M^{k-j})^2\right) \sum_{y \in G} w_y^3,
$$

again using Lemma 1 and noting that each vertex appears in at most $3(M^{k-j})^2$ terms in the sum.

We will now combine these results and use the concentration inequality. We define

$$
\mu_{k,j} = \left(2M^{k-j}\hat{d}\right) \hat{V}_{j-1} \geq \left(2M^{k-j}\tilde{d}\right) V_{j-1} \geq \mu(Y_{k,j})
$$

and note that

$$
\Pr\left[Y_{k,j} > 2\mu_{k,j}\right] \leq \Pr\left[Y_{k,j} > \mu(Y_{k,j}) + \lambda\right],
$$

where $\mu_{k,j} \leq \lambda \leq 2\mu_{k,j}$.

$$
\begin{aligned}
\Pr[Y_{k,j} > 2\mu_{k,j}] &\le exp\left(-\frac{\mu_{k,j}^2}{2(\nu + \frac{a(2\mu_{k,j})}{3})}\right) \\
&\le exp\left(-\frac{\left(2M^{k-j}\hat{d}\right)^2 (\hat{V}_{j-1})^2}{2\left(\left(3(M^{k-j})^2 m\tilde{d}\right) V_{j-1} + \frac{2\left(2M^{k-j}\hat{d}\right)\hat{V}_{j-1}(M^{k-j}m)}{3}\right)}\right) \\
&\le exp\left(-\frac{4\hat{V}_{j-1}}{2\left(3m + \frac{4m)}{3}\right)}\hat{d}\right) \\
&\le exp\left(-\frac{6}{13}\frac{\hat{V}_{j-1}}{m}\hat{d}\right) \\
&= exp(-\Omega(n^\alpha))
\end{aligned}
$$

Thus, $\Pr[Y_{k,j} > 2\mu_{k,j}] \le exp(-\Omega(n^\alpha))$. By the union bound, $Y_{k,j} \le 2\mu_{k,j}$ for all $j \le k \le \ell - 1$ with probability $1 - \ell^2 e^{-\Omega(n^\alpha)} = 1 - e^{-\Omega(n^\alpha)}$.

Proposition 3 *With probability* $1 - exp(-\Omega(n^\alpha))$,

$$V_k \le (4M\hat{d})^k m \qquad \textit{for all} \quad k \in [0, \ell - 1]$$

Proof of Proposition 3: We prove by induction that

$$V_k \le (4M\hat{d})^k m, \tag{6}$$

given that

$$V_{k,j} \le \left(4M^{k-j}\hat{d}\right)\hat{V}_{j-1} \qquad \text{for all} \quad j \le k \le \ell - 1.$$

The result of Proposition 3 will follow immediately, since that event occurs with probability $1 - exp(-\Omega(n^\alpha))$ by Proposition 2.

Equation (6) holds for $k = 0$ since we have $V_0 = Vol(\{u\}) \le m$. Assume now that (6) holds for $[0, k]$ and consider V_{k+1}.

$$
\begin{aligned}
V_{k+1} &\le \sum_{j\in[0,k+1]} Y_{k+1,j} \\
&\le Y_{k+1,0} + \sum_{j\in[1,k+1]} \left(4M^{k+1-j}\hat{d}\right)\hat{V}_{j-1} \\
&\le \Gamma_{k+1}^L(u) + \sum_{j\in[0,k+1]} \left(4M^{k+1-j}\hat{d}\right)(4M\hat{d})^{j-1} m \\
&\le M^{k+1} + M^k m \sum_{j\in[0,k+1]} \left(4\hat{d}\right)^j \\
&\le M^{k+1} m (4\hat{d})^{(k+1)} \\
&= (4M\hat{d})^{k+1} m
\end{aligned}
$$

To obtain the second-to-last line we have assumed that $M \geq 2$.

Proof of Lemma 5:

With probability $1 - e^{-\Omega(n^\alpha)}$,

$$\begin{aligned}
\sum_{k\in[1,\ell]} Vol(\Gamma_{k-1}(u))Vol(N^L_{\ell-k}(v)) &\leq \sum_{k\in[1,\ell]} \left((4M\hat{d})^{k-1}m\right)(2mM^{\ell-k}) \\
&= 2m^2M^{\ell-1}\sum_{k\in[1,\ell]}(4\hat{d})^{k-1} \\
&\leq 4m^2M^{\ell-1}(4\hat{d})^{\ell-1} \\
&= 4m^2(4M\hat{d})^{\ell-1}
\end{aligned}$$

Proof of Theorem 7

Let

$$E = \bigcup_{k\in[1,\ell]} \left(\Gamma_{k-1}(u) \times N^L_{\ell-k}(v)\right).$$

$|C(u,v)|$ is the number of global edges in E. If (u,v) is a surviving long edge, then $|C(u,v)| \geq f$ by Lemma 3. Let E^f denote the set of ordered f-tuples from E with distinct entries. Let B be the event that

$$\sum_{k\in[1,\ell]} Vol(\Gamma_{k-1}(u))Vol(N^L_{\ell-k}(v)) \leq 4m^2(4M\hat{d})^{\ell-1},$$

which occurs with probability $1 - e^{-\Omega(n^\alpha)}$ by Lemma 5.

$$\begin{aligned}
\Pr\left[|C(u,v)| \geq f\right] &\leq \Pr\left[|C(u,v)| \geq f \mid B\right] + \Pr\left[|C(u,v)| \geq f \mid \bar{B}\right] \\
&\leq \Pr\left[|C(u,v)| \geq f \mid B\right] + e^{-\Omega(n^\alpha)}
\end{aligned}$$

We will now bound $\Pr\left[|C(u,v)| \geq f \mid B\right]$. We will first determine E by revealing the sets $\Gamma_0(u) \ldots \Gamma_{\ell-1}(u)$. Critically, Lemma 4 tells us that the potential edges in E are mutually independent and occur with the same probabilities as in G. Thus,

$$
\begin{aligned}
\Pr\left[|C(u,v)| \geq f \mid B\,\right] &\leq \sum_{((x_1,y_1)\ldots(x_f,y_f))\in E^f} \Pr\left[\bigwedge_{i\in[1,f]} (x_i,y_i)\in G\right] \\
&= \sum_{((x_1,y_1)\ldots(x_f,y_f))\in E^f} \prod_{i\in[1,f]} w_{x_i} w_{y_i} \rho \\
&\leq \rho^f \left(\sum_{((x_1,y_1)\ldots(x_f,y_f))\in E^f} \prod_{i\in[1,f]} w_{x_i} w_{y_i} \right) \\
&= \rho^f \left(\sum_{k\in[1,\ell]} Vol(\Gamma_{k-1}(u)) Vol(N^L_{\ell-k}(v)) \right)^f \\
&\leq \rho^f \left(4m^2 (4M\hat{d})^{\ell-1} \right)^f \\
&= \left(4m^2 (4M\hat{d})^{\ell-1} \rho \right)^f
\end{aligned}
$$

Since $\hat{d} = n^\alpha \leq (\frac{nd}{m^2})^{1/\ell} n^{-3/f\ell}$,

$$
\begin{aligned}
\Pr[|C(u,v)| \geq f \mid B] &\leq \left(4m^2 (4Mn^\alpha)^{\ell-1} \rho\right)^f \\
&\leq \left(4m^2 (4M)^{\ell-1} (n^\alpha)^\ell \frac{1}{nd}\right)^f \\
&\leq \left((4M)^\ell n^{-3/f}\right)^f \\
&= O(n^{-3})
\end{aligned}
$$

Thus the probability that a given long edge survives is at most

$$
O(n^{-3}) + e^{-\Omega(n^\alpha)} = O(n^{-3}).
$$

Since there are at most n^2 edges in G, with probability $1 - O(n^{-1})$ no long edges survive. In that case, $L' \setminus L$ contains only short edges, and there are $O(\tilde{d})$ of these by Lemma 2, so part (1) follows. To prove (2), note that if no long edges survive, then all edges in L' must be short. If (u,v) is an edge in L', $d_L(u,v) \leq \ell$. If p' is a path between two vertices x, y in L' with length k, then by replacing each edge with a short path we obtain a path p in L between x and y with length at most ℓk. The result follows.

7.2 Proof of Theorem 8

The $G(n,p)$ model with $p = dn^{-1}$ is a special case of $G(w)$ with $d = \tilde{d} = m$. Our bound from from Lemma 5 becomes

$$
\sum_{k\in[1,\ell]} Vol(\Gamma_{k-1}(u)) Vol(N^L_{\ell-k}(v)) \leq 4(4M)^{\ell-1} d^{\ell+1}.
$$

We then obtain

$$
\begin{aligned}
\Pr\left[|C(u,v)| \geq f \mid B\,\right] &\leq \rho^f (4(4M)^{\ell-1} d^{\ell+1})^f \\
&= (4(4M)^{\ell-1} d^\ell n)^f
\end{aligned}
$$

and the result follows.

7.3 Proof of Theorem 9

We will use the following lower bound on neighborhood size (see [9] p. 260).

Lemma 6 (Neighborhood lower bound)

Let ℓ be a fixed constant, If $d \geq n^{1/\ell}(\log(n^2))^{1/\ell}$, then

$$\Pr\left[|N_{\ell-1}^{G}(x)| < \frac{5}{6}(n\log(n^2))^{1-1/\ell}\right] < n^{-4},$$

provided that n is sufficiently large.

Let $d \geq 6fn^{\frac{1}{\ell}}(\log n)^{\frac{1}{\ell}}$, as in the statement of the theorem. Let u, v be any pair of vertices in H. We will show that G contains f short disjoint paths from u to v, which will imply that every edge in $G \setminus L$ survives. Partition the vertices $V \setminus \{u, v\}$ into f disjoint sets $V_1 \dots V_f$, each of size n/f and let G_i be the induced subgraph of G on $V_i \cup \{u, v\}$. We will ignore the fact that we may not be able to partition into sets of size exactly n/f, since it will not be significant. We can view G_i as a $G(n, p)$ random graph with

$$d' = 6n^{\frac{1}{\ell}}(\log n)^{\frac{1}{\ell}} \geq 6n^{\frac{1}{\ell}}(\log n)^{\frac{1}{\ell}} \geq 6|G_i|^{\frac{1}{\ell}}(\log|G_i|)^{\frac{1}{\ell}}.$$

By applying Lemma 6 to any particular G_i,

$$|N_{\ell-1}^{G_i}(x)| \geq \frac{5}{6}(|G_i|\log(|G_i|^2)^{1-1/\ell}$$

with probability at least $1 - |G_i|^{-4} = 1 - (n/f)^4$. With probability at least $1 - f(n/f)^4 \geq 1 - n^{-4}$ this holds for all $G_1 \dots G_f$, and we let A denote this event. If A holds, there is likely to be an edge in G from $N_{\ell-1}^{G_i}(x)$ to v.

$$\begin{aligned}
\Pr\left[\text{No edge from } N_{\ell-1}^{G_i}(x) \text{ to } v \mid A\right] &\leq (1-p)^{|N_{\ell-1}^{G_i}(x)|} \\
&\leq exp(-p|N_{\ell-1}^{G_i}(x)|) \\
&\leq exp(-p\frac{5}{6}(|G_i|\log(|G_i|^2)^{1-1/\ell}) \\
&\leq exp(-6fn^{-1}n^{\frac{1}{\ell}}(\log n)^{\frac{1}{\ell}}\frac{5}{6}(|G_i|\log(|G_i|^2)^{1-1/\ell}) \\
&\leq exp(-5\log(n/f)) \\
&\leq O(n^{-5})
\end{aligned}$$

Thus, conditional on A, there is an edge from $N_{\ell-1}^{G_i}(x)$ to v for each $i \in [1, f]$ with probability $1 - fO(n^{-5}) = 1 - O(n^{-5})$. The event A occurs with probability $1 - O(n^{-4})$. Thus, with probability $1 - O(n^4)$ there exist f short disjoint paths from u to v, and hence the ℓ-flow from u to v is at least f. Since there are at most n^2 edges in G, by the union bound every edge in $G \setminus L$ survives with probability $1 - O(n^2)$.