HELSINKI UNIVERSITY OF TECHNOLOGY
T-79.300 Postgraduate Course in Theoretical Computer Science
Fitness Landscapes (Spring 2002)

Non-uniform random graphs

Satu Virtanen
satu@cs.hut.fi

Notational conventions

G	a graph $G = (V, E)$
V	the set of vertices of a graph
E	the set of edges of a graph
n	the number of vertices in a graph, $n = \|V\|$
m	the number of edges in a graph, $m = \|E\|$
k	the degree of a vertex
$\bar{k}$	the average degree of a vertex in a graph

1 Introduction

Within the past few years, there has been a rapidly growing interest in realistic modeling of complex systems, such as social networks, chemical structures, neural systems, and communication networks such as the Internet. With even a little knowledge of networks, it is quite evident that these and many other complex systems can be modeled — to some extent — by networks. Many believe that networks are a modeling tool that can represent any complex natural system with sufficient accuracy for meaningful analysis [LM01]. The entire system can rarely be captured in just one simplified network model, but many important aspects can be modeled and studied with ease by the use of network models.

Since the research effort on natural networks has expanded, in part due to the 1998 article in Nature by Watts and Strogatz [WS98], it has become clear that the previously existing models for networks do not capture some of the essential properties inherent in natural networks of interest. Many new models have been proposed and studied during these past years, one improving the other. In order to understand the approaches and results, some graph theoretical background is necessary.

2 Graphs

The mathematical equivalent of a network is called a *graph*. A graph is a pair of sets $G = (V, E)$, where the elements of V are called the *vertices* of the graph (*i.e.* the nodes of the network) and those of E are *edges* (*i.e.* the connections in the network). The cardinalities of the sets V and E are denoted by $|V| = n$ and $|E| = m$. For a comprehensive review on graphs as well as an introduction to random graphs, see [Die00].

An edge connects two vertices to each other. The connection can be *directed* or *undirected*; in a directed graph, an edge $(v, w) \in E$ (where $v, w \in V$) is thought to have a direction "from v to w" (as a "one-way street"). Here the graphs are in general considered to be undirected, *i.e.* (v, w) and (w, v) both represent the same two-way connection. The edges could also be assigned *weights*; in an unweighted graph, all weights are considered to be the same, usually one.

A common way to represent a graph is by its *adjacency matrix* $\mathbf{A}$, which is an $n \times n$ matrix, with one row and column per each vertex v_i, and $a_{ij} = 1$ exactly when $(v_i, v_j) \in E$, zero elsewhere. The *degree* k_i of a vertex v_i is the number of edges joined to v_i, *i.e.* the row sum of the ith row of $\mathbf{A}$ (or respectively the ith column, as for undirected graphs $\mathbf{A}$ is symmetric). If all vertices have the same degree k, the graph is *k-regular*. For a non-regular graph, the *average degree* $\bar{k}$ of the graph is simply $\frac{1}{n} \sum_{i=0}^{n} k_i$.

The important property of a graph, entirely determined by its adjacency matrix $\mathbf{A}$, is its *connection topology*, *i.e.* the manner in which the vertices are connected. The connection topology may be *regular*, *random*, or something in between of these two extremes [WS98]. Examples of regular graphs are lattices such as a ring of vertices all connected to the same number of their nearest neighbor vertices on each side.

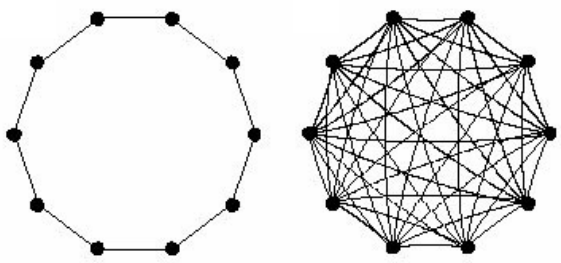

Figure 1: Two ring lattices; with connections only to "nearest neighbors" on the left, fully connected on the right [Str01]

Random graphs have traditionally been generated by first fixing a set of vertices and then choosing edges to be present randomly, using some probability distribution. As Erdős and Rényi constructed such models of random graphs decades ago with little data on large networks, their models could not be examined in

depth by empirical tests. Nowadays these models can be finally compared to natural systems as we now have a plenitude of topological information on huge complex systems available [BA99a].

3 Measures of graph structure

In practice, natural networks can rarely be accurately modeled by either regular or random graphs; something in between is required to capture the essential topologic information. Examples of such networks are neural networks, chemical bond structures, social and computer networks, or food webs [BW00, ASBS00]. There is a need for a certain degree of randomness, and yet uniformity is not desired in generation of such network structures.

Several measures can be calculated for a graph to reflect its degree of regularity or randomness. Among others, the connectivity of a local neighborhood of a vertex (*i.e.* how many edges are there between close-by vertices) has been considered. For a graph of n vertices, there can be at most $n(n-1)/2$ different edges between the vertices if every vertex in connected to every other vertex (a simple combinatorial observation, reflexive edges excluded). For a single vertex v_i and its neighborhood (defined by some measure of "spatial nearness" between vertices), a *clustering coefficient* $\mathcal{C}_i$ can be defined as follows: if v_i has c_i neighbors, there can exist at most $c_i(c_i-1)/2$ edges between v_i and its neighbors $\Gamma(v_i)$; the fraction of these edges that are really present in the graph is the local clustering coefficient $\mathcal{C}_i$ of the vertex v_i. For the entire graph, the clustering coefficient $\mathcal{C}$ is the average over all $\mathcal{C}_i$. [BW00]

Another such quantity is related to the distance between two nodes v_i and v_j calculated as the minimal number of edges $d_{i,j}$ that must be traversed to reach v_j from v_i (or vice versa; the result is the same for an undirected graph). Calculating this for each pair of vertices in G and averaging, the *characteristic path length* $\mathcal{L}$ of G is obtained. This is a meaningful calculation only when the graph is *connected*, *i.e.* any vertex can be reached by a finite number of edge traversals from any other vertex. Otherwise the averaging will have to cope with infinite distances. The connectivity requirement ensures that $d_{i,j}$ is positive and finite $\forall i \neq j$ [LM01].

The characteristic path length is thus a global property of the graph, measuring the typical distance between two vertices, whereas the clustering coefficient is a local property, measuring the clustering (*i.e.* the connectivity of the neighborhood for a vertex). For a regular graph, clustering is high, whereas random graphs are not clustered. When the system size n grows, the characteristic path length $\mathcal{L}$ for many regular graphs grows linearly with respect to n, whereas for a random network the growth is only logarithmic [WS98]. The *proximity ratio* μ of a graph is defined to be the ratio $\mathcal{C}/\mathcal{L}$, normalized by $\mathcal{C}_{\text{rand}}/\mathcal{L}_{\text{rand}}$ *i.e.* the respective figures for a completely random graph [Wal99]. This of course produces $\mu = 1$ for random graphs.

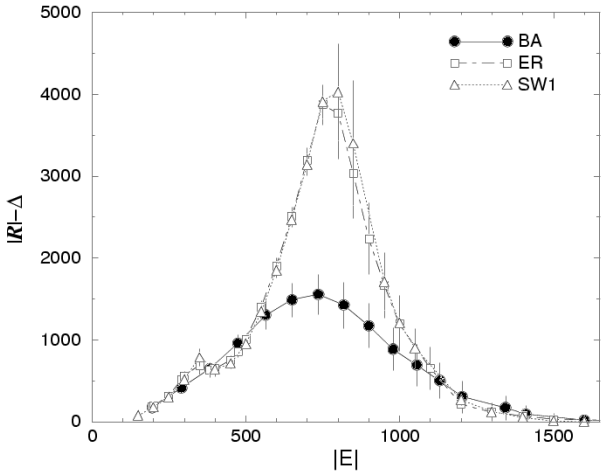

Figure 2: Number of relevant cycles longer than triangles $\mathcal{R} = |R| - \triangle$ for three random graph models to be presented in the next section [GSWF]

Yet another way to characterize the a graph is its *degree distribution*. For a regular graph, all vertices have the same degree, and for a random graph, the distribution depends on the generation method. Also *efficiency* $\mathcal{E}$ of a graph can be useful. It measures the efficiency of information exchange over the graph. In [LM01], efficiency is defined for unweighted, sparse ($m \ll n(n-1)/2$) and connected graphs. The local efficiency of a graph represents how *fault tolerant* the graph is, *i.e.* how efficient the information exchange will be in the neighborhood of vertex v_i if v_i itself is removed. In a sense related to fault tolerancy is also the number of *relevant cycles* $\mathcal{R}$ *i.e.* paths that start and end in the same vertex that cross more than three edges. It estimates both the degree of connectivity and the path length and has been discussed by Gleiss et al. [GSWF]

Hogg [Hog98] has proposed *approximate entropy* as a measure for "similarity between substructures" [Wal99]. It is calculated from the frequences with which different substructures appear in the structure at hand. Approximate entropy essentially measures a similar property of a graph as the clustering coefficient $\mathcal{C}$. Hogg has studied as examples Coloring and k-Sat.

4 Real-world examples

In the table below , examples of the measures introduced in the previous section are given for some real-world networks to serve as an example. The clustering coefficient $\mathcal{C}$ and characteristic path length $\mathcal{L}$ have been normalized by the respective figures on random graphs of same size n and average degree $\bar{k}$. The connectivity lengths $\mathcal{D}$ (both the local and global version) have also been normalized in the same fashion. The second to last column displays the value of γ for those networks that have been found to obey the power law $P(k) = k^{-\gamma}$ and "exp" the one that is known to decay exponentially.

To fit the table on the page width, explanations on the networks and references are explained below separately.

Network	$\mathcal{C}_n$	$\mathcal{L}_n$	μ	$\mathcal{D}_{\text{global}}$	$\mathcal{D}_{\text{local}}$	$\mathcal{E}_{\text{global}}$	$\mathcal{E}_{\text{local}}$	γ
IMDB	2925	1.22	2398	1.12	0.001			2.3 ± 0.1
WSPG	16.00	1.51	10.60					exp
CE	5.60	1.18	4.75	1.01	0.17	0.46	0.47	
MBTS				0.14	0.10	0.63	0.03	
WWW	0.18					0.28	0.36	≈ 2.1
Internet					0.29	0.26		2.48

The IMDB network is the film actors' collaboration network constructed from the International Movie database (`http://www.imdb.com`). $\mathcal{C}$ and $\mathcal{L}$ have been calculated for the largest connected component ($n = 225226$ and $\bar{k} = 6$), but $\mathcal{D}$'s for the entire network [WS98, ML00, Wat99].

The WSPG network represents the the Western States Power Grid and has $n = 4941$ and $\bar{k} = 2.67$ [WS98, Str01, Wat99].

The CE network is the neural network of the nematode worm *Caenorhabditis elegans* whose neural network is completely known, with $n = 282$ and $\bar{k} = 14$ [WS98, ML00, LM01, Wat99].

The MBTS network portrays the Massachusetts Bay transportation system with $n = 124$ and $m = 124$; the physical distances between the subway stations have been taken into account [ML00].

The WWW network represents a portion of the World Wide Web as documented at `http://www.nd.edu/~networks` with $n = 325729$ documents and $m = 1090108$ links (for *clust*, also a value of 0.3 is given), whereas the Internet's representative network is constructed according to the database at `http://moat.nlanr.net` with $n = 6474$ and $m = 12572$ (supposedly the network of domains) [LM01].

5 Generation methods of non-uniform random graphs

There are two common models for generating uniform random graphs by Erdős and Rényi (the ER-model). Familiarity with these models is necessary in order to examine the forthcoming proposals for non-uniform models. In the $G_{n,m}$ model, the generation process starts with initial n vertices and adds m edges randomly, by uniform sampling of the $n(n-1)/2$ possible edges [Die00, Wal01].

Another approach is not to fix the number of edges in advance, but rather the probability for an edge to be included in the graph: the $G_{n,p}$ model again starts with n initial nodes and a fixed probability p for an edge to be included, thereby obtaining $\mathbb{E}[m] = p \cdot n(n-1)/2$ [Wal01].

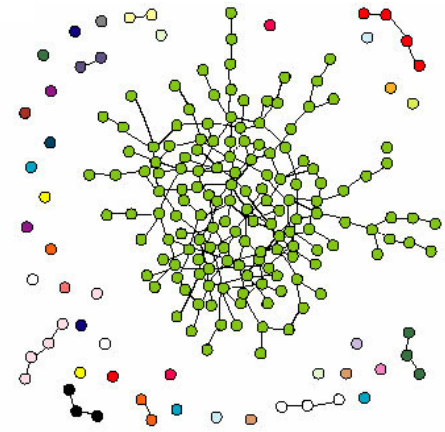

Figure 3: An example of a $G_{n,m}$ graph with $n = 200$ and $m = 193$ [Str01]

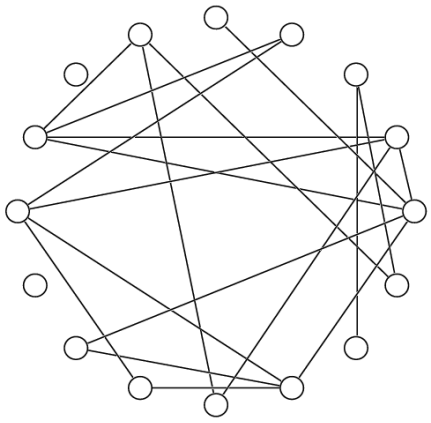

Figure 4: An example of a $G_{n,p}$ graph with $n = 16$ and $p = 1/7$ [NWS02]

For $G_{n,p}$, the degree distribution $P(k)$ (*i.e.* the probability for a vertex to have degree k) is a Poisson distribution $P(k) = e^{-\lambda}\lambda^k/k!$, where λ is $(n-1)p$, the expected node degree — thus vertices that have a high degree are unlikely, due to the exponential decay of the Poisson distribution [Wal01]

5.1 Small-world networks

The class of non-uniform random graphs introduced by Watts and Strogatz [WS98] has become known as *small-world* networks (also the WS-model). The rewiring-procedure for generating such graphs is the following: a ring lattice with n vertices is taken as the initial structure, with each vertex having degree $k = 2h$, h edges on each side connecting the edge to the nearest neighbors both in clockwise and counterclockwise direction.

Each vertex is visited in turn, and each of its k edges is randomly "rewired" with probability p, *i.e.* the target vertex of the edge is changed to some other vertex

with probability p (excluding vertices that are already directly connected to the present vertex, as well as the present vertex itself).

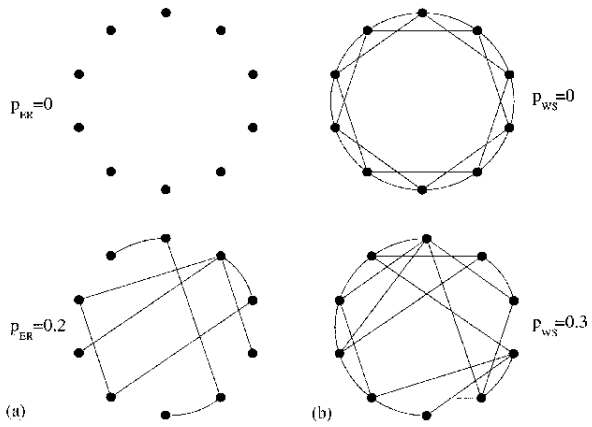

Figure 5: Comparison of ER-graphs (a) and WS-graphs (b) for $p = 0$ and a larger p (0.2 for ER and 0.3 for WS) [BAJ99].

For $p = 0$, the structure remains regular, and for $p = 1$, all edges are randomly rewired at one end. It should be noted that even for $p = 1$, $k_i \geq h$ for all vertices v_i, so the resulting network can not be considered a "true" random graph. A careful mathematical analysis on this observation can be found in [BW00].

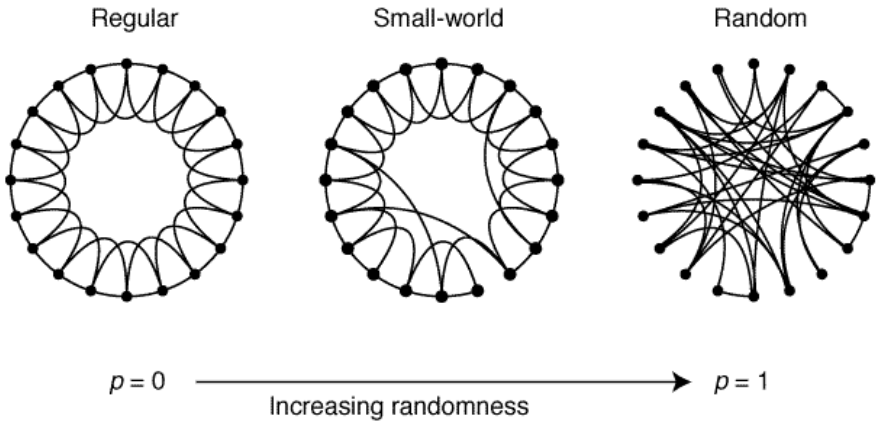

Figure 6: Rewiring a ring lattice ($n = 20$ and $h = 2$) [WS98]

The interesting observation is the behavior of the characteristic path length $\mathcal{L}$ and the clustering coefficient $\mathcal{C}$ when p is varied. In order for $\mathcal{L}$ to be meaningfully defined, the resulting graph must be connected; taking $k \gg \ln n$ should guarantee this [WS98]. Watts and Strogatz [WS98] have observed that over a broad interval, $\mathcal{L}(p) \approx \mathcal{L}_{\text{rand}}$ while $\mathcal{C}(p) \gg \mathcal{C}_{\text{rand}}$. The graphs for which this applies are said to have the small-world property. $\mathcal{L}(p)$ drops rapidly even after a small amount of rewiring, due to the introduction of "shortcuts", *i.e.* edges

that "span" further than any of those in the regular lattice, whereas the small amount of rewiring does not suffice to eliminate the clustering of the regular lattice. Even a small density of such shortcuts brings the typical separation distance close to that of random graphs [NMW99].

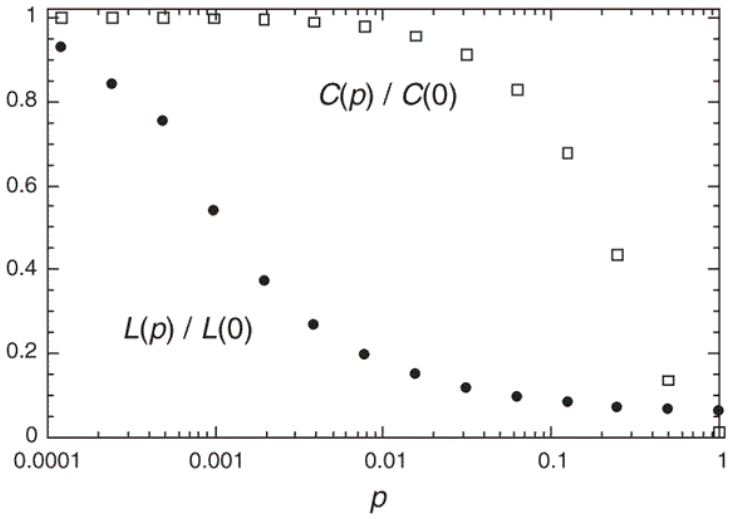

Figure 7: Normalized values of $\mathcal{C}(p)$ and $\mathcal{L}(p)$ for a similar lattice than in Figure 6 [WS98] with $n = 1000$ and $h = 5$

The other measures of the previous section can also be used to characterize the appearance of the small-world property: for small-world graphs $\mu \gg 1$, whereas the proximity ratio is smaller for regular graphs and one for random graphs [Wal99]. Also, graphs that are both locally and globally efficient, possess the small-world property [LM01].

It has been shown that small-world behavior is a crossover phenomenon and not a phase transition [BA99b]. The appearance of the small-world property is dependent on both the network size (it is not visible in small networks) and the rewiring probability p. This matter is also discussed in [DM00], where the goal is to find a scaling form for $\mathcal{L}(p, n)$. Path lengths are also the focus in [NMW99], where an analytic solution for the distribution of the path lengths of a small-world graph is found by mean-field approximation (*i.e.* using average values to represent distributions).

Another interesting observation is that small-world networks can make search problems difficult as shortcuts can "mislead" heuristics. Walsh [Wal99] conjectures that "exceptionally hard problems will be more common in constraint satisfaction problems whose constraint graph has a small world topology than those with a purely random topology".

6 Alternative approaches

One of the disturbing setbacks with the method of random rewiring by Watts and Strogatz is that it cannot be easily analyzed. For example, it is not guaranteed that the process always produces a connected graph. Gent et al. [GHPW99] propose another approach for generating small-world networks. Their most comprehensive example was to start with a two-way grid of "streets" and little by little make some of the streets "one-way" at random. It is not a surprise that this increased the expected length of a route from one place to another randomly selected location. What surprised them at first was that the median cost of finding an optimal route from one location to another dropped. They expect this to be the result of reduction in choices to be made.

6.1 Morphing

Gent et al. call the process used for their experiment *morphing*, which is described through three alternative approaches, each intended for "morphing" different kinds of regular structures into more random ones. All types of morphing begin with two structures S_1 and S_2, the first one being a regular structure, and the second one random (ER-model).

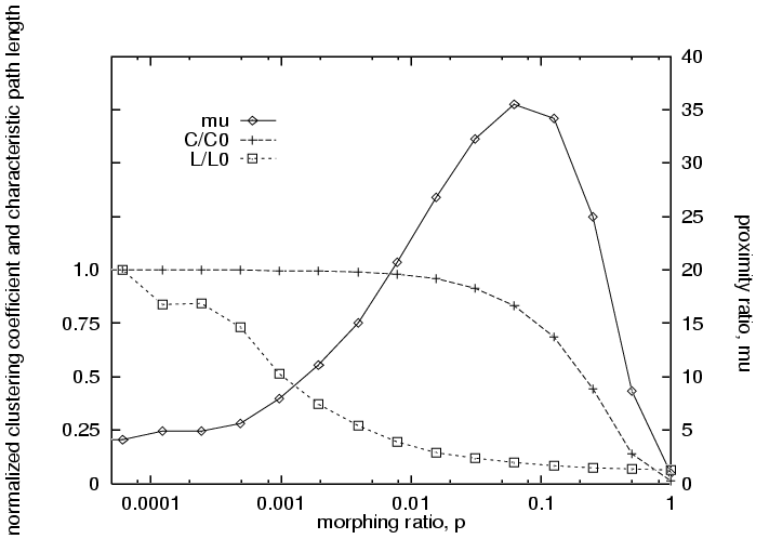

Figure 8: $\mathcal{C}$ and $\mathcal{L}$ for graphs ($n = 1000$, $k = 10$) generated by morphing [GHPW99]; compare with Figure 9.

In the first and second types of morphing, it is required that these structures consist of some kind of substructures that can be separated from the original structures and recombined into a new one. In "type A morph", substructures are taken from S_1 with probability $(1 - p)$ and from S_2 with probability p to form the "randomized" structure (*e.g.*, matrices and SAT instances). In "type B", a fraction $1 - p$ of the new structure is taken from S_1 and p from S_2 (*e.g.*, graphs

and sets). In the third morphing type, "type C", addition and scalar multiplication are expected to be properly defined such that the new structure can be calculated as $(1-p)\cdot S_1 + p \cdot S_2$ (*e.g.*, vectors and functions).

Thus small-world networks are generated by morphing between *e.g.*, a ring lattice or some other regular, clustered structure, and a true random graph. Gent et al. state that "the theoretical analysis of morphs is likely to be much easier than that of rewired graphs" [GHPW99].

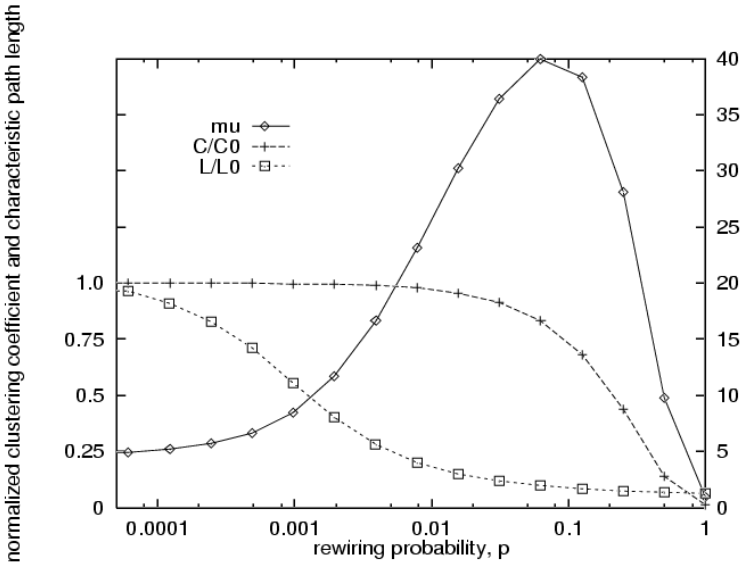

Figure 9: $\mathcal{C}$ and $\mathcal{L}$ for graphs ($n = 1000$, $k = 10$) generated by random rewiring [GHPW99] (same than in Figure 7)

They have experimented also with COLORING, the instances being small-world networks, by encoding the problem into a SAT instance. The conclusion of the work is that "the distribution of search costs across the test set, as well as on individual instances, displays a heavy tail" and that a local search performs poorly on regular structures and well on random ones, whereas a complete search does quite the opposite [GHPW99].

6.2 Metrical graphs

In addition to being "analytically unapproachable", the random rewiring method suffers from heavy limitations set on the graph (undirected, unweighted, connected, sparse etc.), which is why Marchiori and Latora [ML00] propose another approach that is applicable for generic metrical graphs, in principle also disconnected ones. The vertices of the graph are also defined to have a *physical distance* in addition to the distances in the graph $d_{i,j}$, which is also redefined not to be the mere "edge count" but instead the minimal sum of physical distances of all paths connecting vertices v_i and v_j.

Marchiori and Latora propose a measure they call *connectivity length* $\mathcal{D}$ to replace $\mathcal{L}$ and $\mathcal{C}$ used in the WS-model. $\mathcal{D}$ can be computed for all metrical graphs and is said to display the same information as $\mathcal{L}$ and $1/\mathcal{C}$ in characterizing the locality and globality of the network topology (*i.e.* small-world networks have a small $\mathcal{D}$ both locally and globally). The connectivity length $\mathcal{D}$ portrays the efficiency of information propagation (the smaller the better), defined through the separation distances $d_{i,j}$. Thus $\mathcal{D}$ is the amount of separation to which every $d_{i,j}$ should be set in order to maintain the same *performance i.e.* the total amount of information propagated per unit time. The noteworthy point here is that $\mathcal{D}$ of a graph G is not calculated as the arithmetic but the *harmonic mean* of the distances $d_{i,j}$ as the graph is not necessarily connected:

$$\mathcal{D}(G) = H(\{d_{i,j}\}_{i,j \in G}) = \frac{n(n-1)}{\sum_{i,j \in G} 1/d_{i,j}}.$$

6.3 Method of preferential attachment

Another serious drawback of the WS-model is that vertices with a high degree are quite common in real-world graphs, and the WS-model clearly produces a narrow distribution of vertex degrees. Also in ER-graphs, high-degree vertices are rare although the degree distribution is not as narrow as for WS-graphs. The effect that high-degree vertices have on a graph, is that they keep the average path length short due to their heavy connectivity. [Wal01]

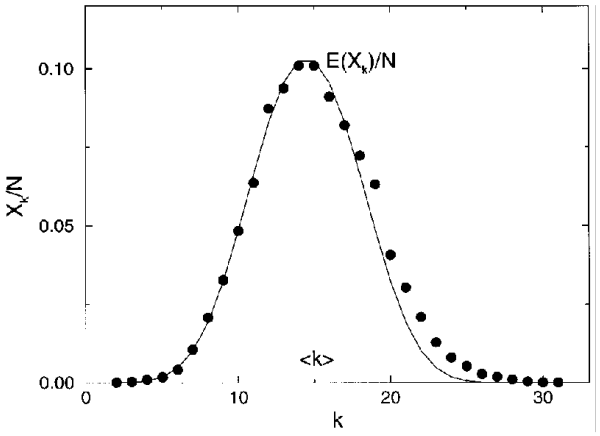

Figure 10: The degree distribution of ER-model (a $G_{n,p}$ graph with $n = 10000$, $p = 0.0015$) [AB02]

A model proposed by Hogg [Hog96] overcomes the degree distribution problem: vertices are grouped into a binary tree and an *ultrametric distance* $u_{i,j}$ is defined for each pair of vertices by calculating the distance up in the tree to the lowest common ancestor. The relative probability of two vertices v_i and v_j being connected by an edge in the graph under construction is $p^{u_{i,j}}$, where p is a fixed probability. For $p = 1$, purely random graphs result from this process, and for

$p < 1$, hierarchial clustering appears. This procedure generates graphs that have more vertices of high and low degrees than in traditional random graphs, but the setback is that they do not have the small-world property, as the clustering is not sufficiently dense [Hog96, Wal99, Wal01].

A power law model proposed by Barabási and Albert (the BA-model) overcomes these difficuleties; it is capable of producing graphs that possess the small-world property and still contain vertices of a high degree [Wal01]. It also avoids another unnatural limitation: the fixed graph size n, is rare in nature but a basic assumption in the models discussed thus far.

Real-world networks are rarely closed systems, on the contrary the number of vertices tends to grow in time so some means of growth should be present in a model that intends to capture the true behavior of real-world systems. Yet another problem is caused by the assumption inherent in both the ER- and the WS-model is that the probability of two vertices being connected were random and uniform, whereas in real-world networks, edges are added by the principle of *preferential attachment*: a high-degree vertex is more likely to "draw" more connections than a more isolated one. [BA99a]

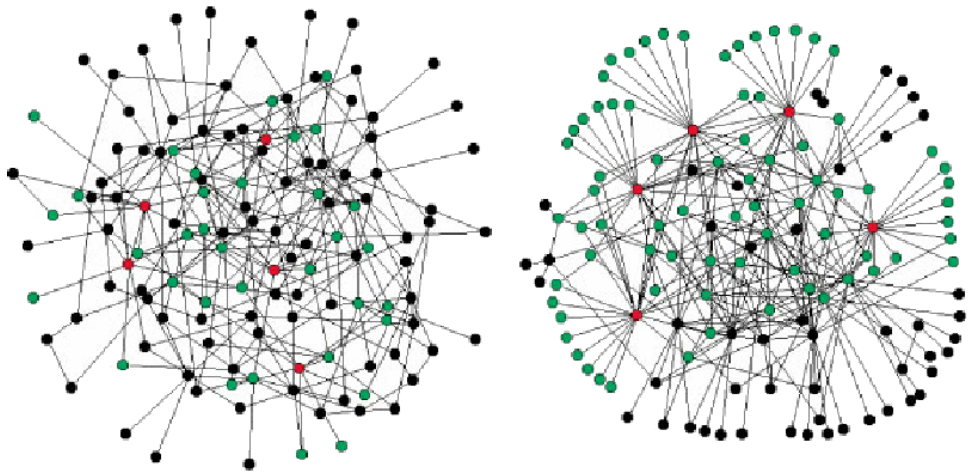

Figure 11: Example of an ER-graph and a BA-graph ($n = 130$, $m = 215$, $\bar{k} = 3.3$) [AJB00]

Barabási and Albert [BA99a] show that the probability that a vertex has degree k decays as a power law, $P(k) \approx k^{-\gamma}$, "independent of the system and the identity of its constituents", which indicates that the graph will eventually be organized into a *scale-free* state. They reach this result by using the following "attachment probability" to attain preferential attachment: as a new vertex v_h is added to the system, the probability that an edge (h, i) is added depends on the magnitude of k_i when compared to the overall degree levels:

$$\Pi(h, i) = \frac{k_i}{\sum_j k_j}.$$

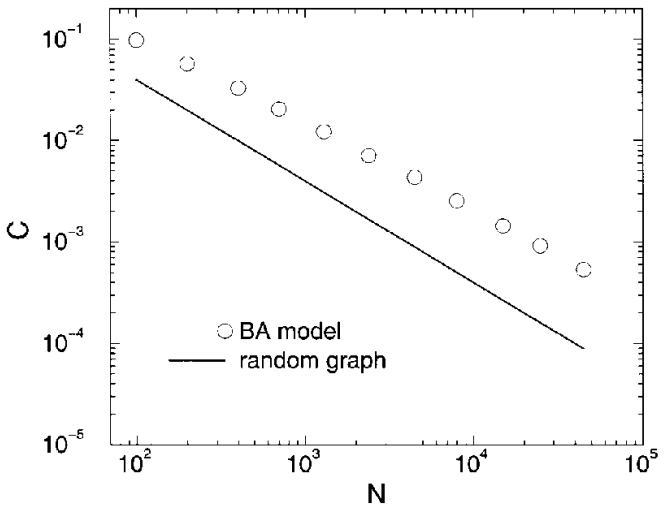

Figure 12: $\mathcal{C}$ for random graphs and graphs generated by combining growth and preferential attachment; $\bar{k} = 4$ [AB02].

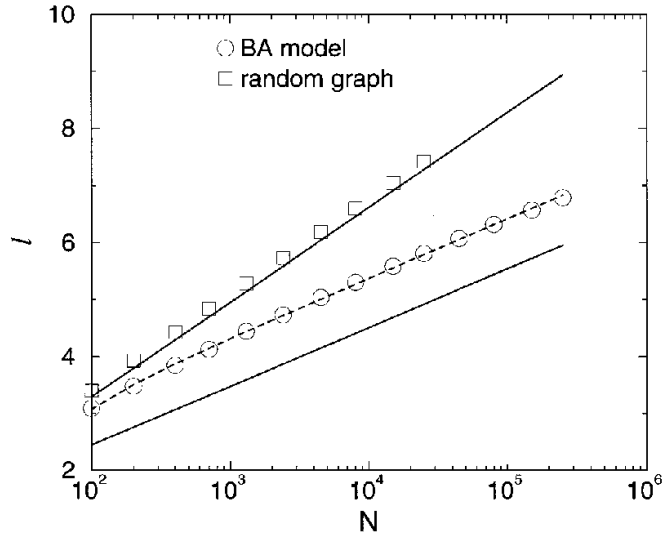

Figure 13: $\mathcal{L}$ for random graphs and graphs generated by combining growth and preferential attachment; $\bar{k} = 4$ [AB02].

Thus the system keeps growing vertex by vertex, exhibiting preferential attachment, and as $P(k)$ is not dependent on the point of time (and therefore not affected by the system size either), self-organizing into a steady, scale-free state. Together with Jeong [BAJ99], they have also showed that both the element of growth and the preferential attachment are necessary in order to produce a scale-free network obeying such a power law, thus behaving like a real-world system.

The BA-model has also be studied by Walsh [Wal01], who studies the cumulative probability of a vertex having *at least* degree k:

$$P(k) = \sum_{i=1}^{k} p(i).$$

He proposes a minor modification to the model in order to avoid having the average degree $\bar{k}$ bounded by the initial size of the graph. Walsh modifies the preferential attachment formula to the following, where c is the number of edges added per each new vertex:

$$\pi(h,i) = \min\left\{1, c \cdot \frac{k_i}{\sum_j k_j}\right\}.$$

Walsh states that this modification is "similar to moving from the $G_{n,m}$ to the $G_{n,p}$ model of random graphs" [Wal01].

References

[AB02] R. Albert and A. Barabasi. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74, 2002.

[AJB00] R. Albert, H. Jeong, and A. L. Barabási. Error and attack tolerance of complex networks. *Nature*, 406:378–382, July 2000.

[ASBS00] L. Amaral, A. Scala, M. Barthelemy, and H. Stanley. Classes of small-world networks. In *Proceedings of the National Academy of Sciences*, volume 97, October 2000.

[BA99a] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[BA99b] Marc Barth and Luís A. Nunes Amaral. Small-world networks: Evidence for a crossover picture. *Physical Review Letters*, 82:3180–3183, 1999.

[BAJ99] Albert-László Barabási, Réka Albert, and Hawoong Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173–187, 1999.

[BW00] A. Barrat and M. Weigt. On the properties of small-world network models. *Eur. Phys. J.*, B 13:547–560, 2000.

[Die00] Reinhard Diestel. *Graph Theory.* Springer-Verlag, New York, second edition, 2000.

[DM00] S. N. Dorogovtsev and J. F. F. Mendes. Exactly solvable small-world network. *Europhysics Letters*, 50(1):1–7, 2000.

[GHPW99] Ian P. Gent, Holger H. Hoos, Patrick Prosser, and Toby Walsh. Morphing: Combining structure and randomness. In *AAAI/IAAI*, pages 654–660, 1999.

[GSWF] Petra M. Gleiss, Peter F. Stadler, Andreas Wagner, and David A. Fell. Small cycles in small worlds. SFI Working Paper 00-10-058.

[Hog96] Tad Hogg. Refining the phase transitions in combinatorial search. *Artificial Intelligence*, pages 127–154, 1996.

[Hog98] Tad Hogg. Which search problems are random? In *Proc. of AAAI98*, pages 438–443, Menlo Park, CA, 1998. AAAI Press.

[LM01] Vito Latora and Massimo Marchiori. Efficient behavior of small-world networks. *Physical Review Letters*, 87(19), November 2001.

[ML00] Massimo Marchiori and Vito Latora. Harmony in the small-world, 2000.

[NMW99] M. Newman, C. Moore, and D. Watts. Mean-field solution of the small-world network model. SFI Working Paper 99-09-066, 1999.

[NWS02] M.E.J. Newman, D.J. Watts, and S.H. Strogatz. Random graph models of social networks. In *Proc. Natl. Acad. Sci. USA 99*, pages 2566–2572, 2002.

[Str01] Steven H. Strogatz. Exploring complex networks. *Nature*, 410:268–276, 2001.

[Wal99] Toby Walsh. Search in a small world. In Dean Thomas, editor, *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99-Vol2)*, pages 1172–1177, S.F., July 31–August 6 1999. Morgan Kaufmann Publishers.

[Wal01] Toby Walsh. Search on high degree graphs. In *IJCAI*, pages 266–274, 2001.

[Wat99] Duncan J. Watts. *Small Worlds*. Princeton University Press, 1999.

[WS98] D. Watts and S. Strogatz. Collective dynamics of 'smallworld' networks. *Nature*, (393):440–442, June 1998.