ELECTRONIC RESEARCH ANNOUNCEMENTS
OF THE AMERICAN MATHEMATICAL SOCIETY
Volume 9, Pages 152–161 (December 18, 2003)
S 1079-6762(03)00123-9

A STOCHASTIC COMPLEX NETWORK MODEL

DAVID J. ALDOUS

(Communicated by Ronald L. Graham)

ABSTRACT. We introduce a stochastic model for complex networks possessing three qualitative features: power-law degree distributions, local clustering, and slowly growing diameter. The model is mathematically natural, permits a wide variety of explicit calculations, has the desired three qualitative features, and fits the complete range of degree scaling exponents and clustering parameters.

1. Introduction

Three popular science books [5], [7], [15], a dozen articles in *Science* and *Nature*, and 154 preprints at `xxx.arXiv.org/cond-mat` deal with *complex networks*, which in this context means the empirical and theoretical study of large graphs, focusing in particular on those possessing the following three qualitative properties, asserted to hold in many interesting real-world examples:

- the degree distribution has power-law tail;
- local clustering of edges: graph is not locally tree-like;
- small diameter — $O(\log(\text{number of vertices}))$.

The nature of that subject—typically not presented as rigorous mathematics—is most easily seen from the long survey papers [1], [8], [13] and associated monograph [12], and we will not attempt an overview here. A shorter survey [6] emphasizes rigorous mathematical results, more of which are appearing in a new journal *Internet Mathematics*. The probability models which have been studied are mostly variations of two basic ideas. In *proportional attachment* models, vertices arrive and form edges to existing vertices with relative probabilities increasing linearly with degree of the existing vertex. In *lattice small worlds* models, vertices are points of $\mathbb{Z}^d$ and in addition to lattice edges there are long edges between random vertex-pairs chosen with probabilities decreasing with distance. But it has often been noted that neither class of models gives all three desired properties without adding extra ingredients to force a property to hold by fiat, as illustrated by the following quotes.

> Most models proposed to describe the topology of complex networks have difficulty capturing simultaneously these two features [in our list of three]. [14]

Received by the editors July 22, 2003.
2000 *Mathematics Subject Classification.* Primary 60K35; Secondary 05C80, 90B15, 94C15.
Key words and phrases. Complex network, Poisson process, PWIT, random graph, scale-free, small worlds, Yule process.
The author was supported in part by NSF Grant DMS-0203062.

> There is a great need for models of real-life networks that incorporate many of the important features of these systems, but still can be analyzed rigorously. [6]

Our purpose is to introduce a model, based on somewhat more sophisticated mathematics, designed to possess all three qualitative properties in a mathematically natural way, and to permit a wide variety of explicit calculations in the n (= number of vertices) $\to \infty$ limit. We do not claim the model is canonical—we could imagine quite different models satisfying the same desiderata—but such models have not appeared in the literature.

A longer account of this model, aimed at non-mathematicians and emphasizing the conceptual aspects of the model and the explicit calculations, will appear elsewhere [3]. In this account, aimed at mathematicians, we emphasize how some existing mathematical structures and methodology can be used to implement the kind of graph-modeling ideas studied more naively in the complex networks literature. The bald construction (Section 2) will at first sight seem neither natural nor tractable. Section 2.1 compares it with standard models. Section 3 reviews some existing mathematical structures and methodology (the Yule process; steady-state Poisson universe; local weak convergence and the PWIT). Section 4 puts it all together to reveal why one can do explicit calculations in the $n \to \infty$ limit, and then Section 4.1 outlines illustrative calculations.

Our model has of course some disadvantages, detailed in [3]. In particular, the graph $\mathcal{G}_n$ is typically not connected. By analogy with classical Erdös-Rényi random graph theory, we expect that (on one side of a critical line in parameter space) there is a unique "giant component" of $\mathcal{G}_n$, whose diameter is $O(\log n)$. One could make a connected graph by modifying the definition so that an arriving vertex always links to its closest neighbor, though this would make explicit formulas more complicated.

2. The finite n model

We define the model in steps (a)–(e).

(a) Vertices $n = 1, 2, 3, \ldots$ arrive successively, vertex n arriving at time $t_n = \log n$. The model can be indexed by number of vertices ($n = 1, 2, 3, \ldots$) or equivalently by time ($0 \le t < \infty$).

(b) When vertex n arrives at time t_n, there are created links from n to each vertex $1 \le j < n$. Each link-length $\bar{d}(n, j, t_n)$ is random with exponential (mean n) law, independent of other link-lengths.

(c) The link-lengths increase with time at deterministic rate 1; at time $t > t_n$ the link (n, j) has length $\bar{d}(n, j, t) = \bar{d}(n, j, t_n)e^{t-t_n}$.

(d) At time t the *distance* $d(i, j, t)$ between two vertices i, j (where $i, j \le e^t$ = number of vertices present) is defined as the minimum, over all paths from i to j, of the length of the path (i.e. sum of link-lengths along the path).

What we have described so far is a "geometry" (we write $\mathcal{D}_n$ or $\mathcal{D}_t$) of an evolving random discrete metric space. The object of this paper is a directed graph process (we write $\mathcal{G}_n$ or $\mathcal{G}_t$) built on top of this geometry. The model has two parameters $\alpha, \lambda > 0$, which enter only via the function

$$p(x) = \min(1, \alpha\lambda e^{-\lambda x}), \quad 0 \le x < \infty. \tag{2.1}$$

We require

$$\int_0^\infty p(x)\ dx < 1, \tag{2.2}$$

which forces the parameter range to be

$$\alpha < 1 \text{ and } \alpha\lambda \leq 1; \text{ or } \alpha\lambda \geq 1 \text{ and } \lambda^{-1} + \log(\alpha\lambda) < 1.$$

In fact the model makes sense for any function $p(x) \leq 1$ satisfying (2.2), but the slightly arbitrary parametrization (2.1) is convenient for obtaining explicit formulas.

(e) When vertex n arrives at time t_n, for each $1 \leq i < n$:
- a directed edge $n \to i$ is created with probability $p(\bar{d}(n, i, t_n))$;
- for each $j < n$ such that $i \to j$ is an existing edge in $\mathcal{G}_n$, a directed edge $n \to j$ is created with probability $p(\bar{d}(n, i, t_n))$.

This provides an inductive description of $\mathcal{G}_n$. Of course as well as its graph structure, $\mathcal{G}_n$ inherits "metric" structure such as real-valued edge lengths (growing exponentially with time).

Let us make two remarks on definition (e). Note that the probabilities depend on link-lengths $\bar{d}(\cdot)$ instead of metric distance $d(\cdot)$. And note that the probability that an existing edge $i \to j$ is "copied" to a new edge $n \to j$ depends on the link-length from n to the "head" i, not to the "tail" j.

2.1. **Comparisons with other models.** In the *small worlds* models, vertices are points in d-dimensional space, which automatically provides a metric distance between vertices, and the model uses some rule to create a random graph with short-range and long-range edges. In purely graph-theoretical models, such as the basic *proportional attachment* model, the vertices have no "intrinsic structure" other than that provided by the graph; we visualize this as saying that each pair of vertices is metric distance 1 apart. In related "copying" models a new vertex picks existing vertices i and, in some random way, creates edges to those vertices i and copies edges to some j for which $i \to j$ is an existing edge, in the spirit of (e) above. One can regard (e) as defining a new class of *metric copying* models, which could be combined with any model of metric distance $d(\cdot)$. In a metric copying model we visualize vertices as points in some abstract metric space, representing (in the case of web pages, say) the difference between the content of the pages, or (for people) some notion of "social distance" based on location, education, profession, interests, etc., of the individuals. In detail our particular geometry model is chosen for mathematical tractability rather than any claimed realism. But it does have the property (see Section 3.3) that the number of vertices within metric distance r of a typical vertex grows exponentially with r. This property is intermediate between, and surely in many contexts more plausible than, the alternatives implicit in the two standard categories of models above.

Of course ours is not the first model seeking to combine the "metric" and "copying" ingredients, but other models [10], [11], [9] do not permit the same range of explicit calculations.

3. Three mathematical structures

Here we recall three structures which will be used in the construction and analysis of our model. A note at the end of each subsection says where they are used.

3.1. **The Yule process.** Fix $0 < \theta < \infty$ and set $N(0) = 1$. The *Yule process* of rate θ, written $(N(t), t \geq 0)$, represents population size in a population of immortal individuals, each of whom reproduces at stochastic rate θ. In other words, it is the continuous-time Markov chain which changes only by $+1$ steps and for which

$$P(N(t+dt) = m+1 | N(t) = m) = \theta m \ dt.$$

This is a well-understood elementary example of a stochastic process, for which we quote two explicit results (both going back to the original 1924 paper by Yule [16]). First

$$N(t) \text{ has geometric}(e^{-\theta t}) \text{ law,} \tag{3.1}$$

meaning $P(N(t) = m) = (1-p)^{m-1}p, \ m \geq 1$ for $p = e^{-\theta t}$. Second, if T is an independent exponential (mean 1) r.v., then

$$P(N(T) - 1 \geq d) = \frac{\Gamma(d+1)\Gamma(1/\theta)}{\theta \Gamma(d+1+\frac{1}{\theta})}, \quad d \geq 0. \tag{3.2}$$

Where used? The Yule process is used in analysing in-degree distribution (Section 4.1.1) and density of induced subgraphs (Section 4.1.3).

3.2. **The steady-state Poisson universe.** In the 1950s the astronomer Fred Hoyle noted that the fact that more distant galaxies recede more rapidly was not logically inconsistent with the supposition that the universe has existed, and looked statistically similar, for ever; because it might be that matter is continually created from empty space. This *steady state* cosmological theory was never widely accepted, but does correspond to a well-defined mathematical model for random points. The simplest "static" model for random points in $\mathbb{R}^d$ is the *Poisson process* of spatial rate 1, for which the mean number of points in a region equals the d-dimensional volume of the region. One now makes a "dynamic" or *space-time* model over time $-\infty < t < \infty$ by specifying that

(i) points move away from the origin as deterministic motion with exponential rate $1/d$; a point at position x at time t will be at position $xe^{(t'-t)/d}$ at times $t' > t$.
(ii) New points appear as a rate-1 space-time Poisson process; that is, the chance of a point arriving in a cube of volume dx during a time interval dt equals $dx\, dt$.

It is straightforward that such a process exists which is time-stationary: at any time the spatial distribution of points is the Poisson process of spatial rate 1. One could also consider this space-time process with the direction of time reversed. Then in (i) the exponential motion is towards the origin; in (ii), a point existing at time t_0 has a "disappearance" time $T < t_0$ such that $t_0 - T$ has exponential (mean 1) law.

Where used? Our model is designed so that its $n \to \infty$ limit is an infinite-dimensional analog of the space-time process above—see paragraph (i) of Section 4.

3.3. **Local weak convergence and the PWIT.** The ideas here are treated in detail in the survey paper [4]. Consider the space of graphs whose edges have real-valued lengths, with one vertex distinguished as the root, and which are *locally finite* in that there are only finitely many vertices within any fixed distance from the root. There is a natural notion of *local convergence* of such graphs. This induces a notion of *local weak convergence* of finite random graphs to a limit infinite (locally finite) rooted graph, based on choosing a uniform random vertex of the finite graph to be a root. In our "geometry" $\mathcal{D}_n$, at time t_n all the $\binom{n}{2}$ link-lengths between the n vertices are independent with exponential (mean n) distribution. It is known (and easy) that $\mathcal{D}_n$ has a $n \to \infty$ limit in the sense of local weak convergence, called the *PWIT*.

The PWIT is defined by a construction, illustrated in Figure 1. Start with a single root vertex $\emptyset$. This root vertex is then given an *infinite* number of links to its *children* vertices, these link-lengths having the law of a Poisson process $(\xi_i^{\emptyset} : 1 \leq i < \infty)$ of rate 1 on $(0, \infty)$. Now, recursively, each vertex v arising as a child of a previous vertex is given an infinite number of links to its children, and these link-lengths $(\xi_i^v : 1 \leq i < \infty)$ again have the law of a Poisson process of rate 1, independent of previous lengths. This procedure is then continued *ad infinitum*. The resulting rooted infinite tree is a well-defined random object, called the *Poisson weighted infinite tree* (PWIT). An obvious, but very useful for calculations, property of the PWIT is its *recursive self-similarity*: the subtrees at each child of the root are independent copies of the PWIT itself.

The *distance* $d(v, w)$ between two vertices of the PWIT is just the sum of link-lengths along the path from v to w. Though we have drawn a tree in Figure 1, the lines merely indicate the link relationships; it is better to think of the edges as absent while retaining the distances $d(v, w)$. In this way we may regard the vertices of the PWIT as an infinite-dimensional analog of the d-dimensional Poisson point process in Section 3.2.

The Yule process is encoded within the PWIT as the process, indexed by $r \geq 0$, which counts the number of vertices within distance r of the root. This is not our main use of the Yule process, but does imply (using (3.1)) that the mean number of such vertices equals e^r, thus providing one justification for regarding the PWIT as "infinite-dimensional".

Where used? Our model adds time-dependence and graph structure to the geometry $\mathcal{D}_t$; we will study its limit behavior in the next section by adding corresponding structure to the PWIT.

In contrast to its use in this paper, the primary use of the PWIT is as a platform to study properties of solutions of combinatorial optimization problems such as the traveling salesman problem on random instances; such study is often intractable in the Poisson model on $\mathbb{R}^d$. See [2] for recent work.

4. The $n = \infty$ limit graph process

We can now explain, in stages, the motivation for our complex network model; it is designed to have tractable structure in the $n \to \infty$ limit.

(i) Just as the $t_0 \to \infty$ limit of the "static" geometry $\mathcal{D}_{t_0}$ is the PWIT, say $\mathcal{D}_0^*$, we can consider the $t_0 \to \infty$ limit of the "dynamic" space-time geometry $(\mathcal{D}_{t_0+t}, -t_0 \leq t < \infty)$, and this limit is a space-time version $(\mathcal{D}_t^*, -\infty < t < \infty)$ of the PWIT. Loosely, we regard this as an infinite-dimensional analog of the steady-state Poisson universe in Section 3.2. At t increases, inter-point distances increase exponentially, but the arrival of new points maintains a constant density and gives the time-stationary property.

(ii) Our finite random graph process $\mathcal{G}_n$ was built on top of $\mathcal{D}_n$ by rule (e) of Section 2. We now use exactly the same rule to define a random graph process

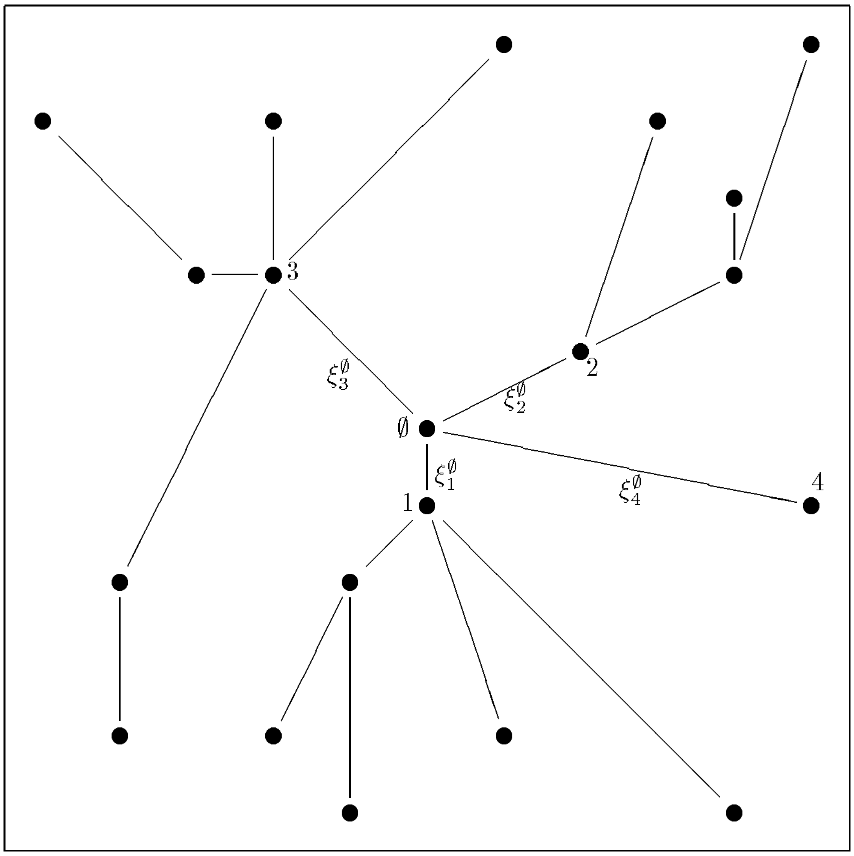

FIGURE 1. **The PWIT.** Illustration of the vertices of the PWIT within a window of radius 3 centered on the root $\emptyset$. Lines indicate the link relationship, but are drawn only when both end-vertices are within the window. Thus the four links of $\emptyset$ shown have lengths $0 < \xi_1^\emptyset < \xi_2^\emptyset < \xi_3^\emptyset < \xi_4^\emptyset < 3$, while there are infinitely many links at $\emptyset$ with lengths greater than 3. Orientation of lines in the figure is arbitrary.

$(\mathcal{G}_t^*,\ -\infty < t < \infty)$ built over the space-time PWIT $(\mathcal{D}_t^*,\ -\infty < t < \infty)$. The time-stationary property of the space-time PWIT extends to $(\mathcal{G}_t^*)$.

(iii) By definition of local weak convergence, the structure of $\mathcal{G}_0^*$ around the root is the $n \to \infty$ limit of the structure of $\mathcal{G}_n$ around a uniform random vertex. So for many quantitative questions about $\mathcal{G}_n$, the $n \to \infty$ limit is given exactly as the answer to a corresponding question about $\mathcal{G}_0^*$. As illustrated later, many such questions can be answered explicitly by doing "calculations around the root" exploiting the structural properties listed below.

While there are many details involved in making the above assertions rigorous, there is nothing mathematically hard. Showing that there exists a *time-stationary* version of $(\mathcal{G}^*_t)$ ultimately comes down to the fact (4.3) that the mean vertex degree is finite. The argument for that fact can be rephrased to show that, starting from no edges and running the dynamics (e), the mean vertex degree remains bounded, so that by taking subsequences of time-averages and passing to a limit one gets a stationary distribution.

Of course, the object $\mathcal{G}^*_0$ that we need to study is not defined explicitly but rather as the stationary distribution of the process $(\mathcal{G}^*_t)$ built over the geometry $(\mathcal{D}^*_t)$. What makes calculations feasible is the interplay of three nice structural properties. We have already mentioned

- recursive self-similarity of the PWIT;
- time-stationarity of the joint process $(\mathcal{D}^*_t, \mathcal{G}^*_t)$;

so let us now develop the third property

- embedded Yule processes.

The basic instance concerns vertex in-degrees. In particular, consider the process

$$N_{\mathrm{in}}(t) = 1 + \text{ in-degree of root at time } T_{\mathrm{arr}} + t,$$

where T_{arr} is the arrival time of the root. This process is the Yule process of rate

$$\beta = \int_0^\infty p(r)\ dr < 1. \tag{4.1}$$

To see why, note that the dynamics of the geometry $(\mathcal{D}^*_t)$ can be expressed intuitively as: for a vertex v present at time t,

$$\begin{array}{c}\text{there is chance } 1 \cdot dr\ dt \text{ that during } [t, t+dt] \text{ a new vertex } v' \\ \text{will appear with a link to } v \text{ of length} \in [r, r+dr]\ .\end{array} \tag{4.2}$$

Apply this at time t to the $N(t) = n$ vertices v which are either the root or have an edge $v \to$ root; an arriving v' has chance $p(r)$ to create a new edge into the root, so the overall rate of such new edges equals $n \int_0^\infty\ p(r)dr$.

4.1. **Illustrative calculations.** Here we outline four calculations from [3], emphasising how the structure is used and omitting calculus details.

4.1.1. *Distribution of in-degree.* The distribution of in-degrees in $\mathcal{G}_n$ has a $n \to \infty$ limit, say D_{in}, which is the in-degree of the root of $\mathcal{G}^*_0$. The age $-T_{\mathrm{arr}}$ of the root of $\mathcal{G}^*_0$ has exponential (mean 1) law, as for the steady-state Poisson process. So D_{in} has the law of $N_{\mathrm{in}}(-T_{\mathrm{arr}}) - 1$, where $N_{\mathrm{in}}(\cdot)$ is the Yule process of rate β at (4.1). So from (3.2) we can read off the formula

$$P(D_{\mathrm{in}} \geq d) = \frac{\Gamma(d+1)\Gamma(1/\beta)}{\beta\Gamma(d+1+\frac{1}{\beta})}, \quad d \geq 0.$$

This has mean

$$ED_{\mathrm{in}} = \tfrac{\beta}{1-\beta} \tag{4.3}$$

and power-law tail

$$P(D_{\mathrm{in}} = d) \sim \beta^{-2}\Gamma(1/\beta)\ d^{-1-\frac{1}{\beta}}.$$

4.1.2. *Distribution of out-degree.* The distribution of out-degrees in $\mathcal{G}_n$ has a $n \to \infty$ limit, say D_{out}, which is the out-degree of the root of $\mathcal{G}_0^*$. Since the out-edges are formed on arrival, we may suppose the root of the PWIT has just arrived. Consider a link (root, v') of length r. For each out-edge of v', and for v' itself, there is chance $p(r)$ that a corresponding out-edge is created at the root, so the number of such edges has Binomial$(1+D^{(i)}, p(r))$ law, where $D^{(i)}$ is the out-degree of v'. Summing over all links whose lengths $(\xi_i, i \geq 1)$ form a Poisson process of rate 1,

$$D_{\text{out}} = \sum_{i=1}^{\infty} \text{Binomial}(1 + D^{(i)}, p(\xi_i)). \tag{4.4}$$

By recursive self-similarity, the $(D^{(i)})$ are independent with law D_{out}, so (4.4) provides an equation for the law of D_{out}. Of course $ED_{\text{out}} = ED_{\text{in}}$ but the distributions are different. While we do not have a useful explicit solution in general, one can calculate moments and can identify the law in several parameter-limit cases [3].

4.1.3. *Density of induced complete subgraphs.* Let G be a finite directed acyclic graph. Define a random variable $\mathbf{X}_n(G)$ to be the number of vertex-subsets V of $\mathcal{G}_n$ such that $\mathcal{G}_n$ restricted to V is isomorphic to G. Our methodology enables us to study asymptotics in terms of $\mathcal{G}_0^*$. Essentially (omitting fussy conventions about rooting and automorphisms of G)

$$\lim_n \tfrac{1}{n} E\mathbf{X}_n(G) = \chi(G)$$

where

> $\chi(G)$ is the expected number of vertex-subsets V of $\mathcal{G}_\infty^*$ including the root such that $\mathcal{G}_\infty^*$ restricted to V is isomorphic to G.

One can obtain formulas for $\chi(G)$ in many cases, in particular for the complete directed acyclic graph K_r on $r \geq 2$ vertices. The idea in this case is that a K_{r+1} is formed by a newly arriving vertex creating edges to all vertices of an existing K_r. For a given K_r, the process which counts the number of such K_{r+1}'s, plus the K_r itself, turns out to be a Yule process of rate $\beta_r = \int_0^\infty p^r(u)\, du$, and we eventually obtain the formula

$$\chi(K_r) = \prod_{u=1}^{r-1} \frac{\beta_u}{1-\beta_u}.$$

The case $r = 3$ is useful for the following reason. One of the basic summary statistics in empirical study of complex networks is the *cluster coefficient*, which measures relative density of triangles. For our purposes we use

$$\begin{aligned} \kappa_{\text{cluster}} &= \lim_n \frac{E(\text{number of triangles in } \mathcal{G}_n)}{E(\text{number of directed 2-paths in } \mathcal{G}_n)} \\ &= \frac{\chi(K_3)}{(ED_{\text{in}})(ED_{\text{out}})}. \end{aligned}$$

So we have formulas for two basic parameters (mean degree and κ_{cluster}) in terms of $p(\cdot)$. As shown in [3], an advantage of our particular specification of $p(\cdot)$ in terms of two parameters α, λ is that the complete range $0 < ED_{\text{in}} < \infty$, $0 < \kappa_{\text{cluster}} < 1$ is attained, for unique parameter values.

4.1.4. *Distribution of edge-lengths.* In our graph model, edges have real-valued lengths, and so there is a law L for the length of a "typical" edge in the limit $\mathcal{G}_0^*$, representing the $n \to \infty$ limit law of a uniform random edge of $\mathcal{G}_n$. We can calculate this law using the time-dynamics of the limit graph process $(\mathcal{G}_t^*)$. We treat the simpler case $\alpha\lambda < 1$. Consider the lengths of the in-edges at a particular vertex v_0. Following a tradition in mathematical probability, we visualize an in-edge of length ℓ as a "particle" at position ℓ on a line; we also put a particle at position 0 to represent the vertex v_0 itself. If we start time τ with $\tau = 0$ at the arrival time of v_0, then the evolution of the "particle process" can be specified as follows.

(i) There is a particle at position 0 at all times $\tau \geq 0$.

(ii) For each particle (at position x at time τ, say), at stochastic rate α per unit time a new particle appears at position $x + \mathrm{Exp}(\lambda)$, where $\mathrm{Exp}(\lambda)$ denotes a r.v. with exponential (rate λ) law.

(iii) Particle positions increase deterministically at exponential(1) rate: a particle at x at time τ will be at $xe^{\tau_0 - \tau}$ at time $\tau_0 > \tau$.

Rule (ii) derives from (4.2): for an existing edge (v', v_0), a new vertex arriving at link-distance r from v' creates an edge to v_0 with probability $p(r)$, so the rate at which each existing edge is copied equals $\int_0^\infty p(x)\, dx = \alpha$; moreover conditional on copying, the distance r has $\mathrm{Exp}(\lambda)$ distribution, and so the length of the new edge equals the length of the old edge $+\mathrm{Exp}(\lambda)$.

Such a particle process can be analyzed by classical methods—set up and solve a differential equation for

$$g(\tau, x) = \text{ mean density (in } x\text{) of particles at time } \tau.$$

Ultimately one obtains a hypergeometric-type series expression for the probability density function of L:

$$\frac{1-\alpha}{\alpha} \sum_{i=0}^{\infty} \frac{(i+1)\Gamma(\alpha+3)\ (-\lambda x)^i}{\Gamma(i+\alpha+3)}, \quad 0 < x < \infty. \tag{4.5}$$

References

[1] R. Albert and A.-L. Barabási, *Statistical mechanics of complex networks*, Rev. Mod. Phys. **74** (2002), 47–97. MR **2003d**:82055

[2] D.J. Aldous and A.G. Percus, *Scaling and universality in continuous length combinatorial optimization*, Proc. Natl. Acad. Sci. USA **100** (2003), 11211–11215.

[3] D.J. Aldous, *A tractable complex network model based on the stochastic mean-field model of distance*, arXiv:cond-mat/0304701, 2003.

[4] D.J. Aldous and J.M. Steele, *The objective method: Probabilistic combinatorial optimization and local weak convergence*, in Probability on Discrete Structures, H. Kesten (ed.), 1–72, Springer, 2003.

[5] A.L. Barabási, *Linked: the new science of networks*, Perseus Press, Cambridge, MA, 2002.

[6] B. Bollobás and O. Riordan, *Mathematical results on scale-free random graphs*, Handbook of Graphs and Networks (S. Bornholdt and H.G. Schuster, eds.), Wiley, 2002.

[7] M. Buchanan, *Nexus: Small worlds and the groundbreaking science of networks*, W.W. Norton, 2002.

[8] S.N. Dorogovtsev and J.F.F. Mendes, *Evolution of networks*, Adv. Phys. **51** (2002), 1079–1187.

[9] A. Fabrikant, E. Koutsoupias, and C.H. Papadimitriou, *Heuristically optimized trade-offs: a new paradigm for power laws in the internet*, International Colloq. Automata, Languages and Programming, 2002.

[10] J. Jost and M.P. Joy, *Evolving networks with distance preferences*, Physical Review E **66** (2002), 036126. MR **2002i**:37133

[11] F. Menczer, *Growing and navigating the small world web by local content*, Proc. Natl. Acad. Sci. USA **99** (2002), 14014–14019.

[12] J.F.F. Mendes and S.N. Dorogovtsev, *Evolution of networks: From biological nets to the internet and WWW*, Oxford Univ. Press, 2003.

[13] M.E.J. Newman, *The structure and function of complex networks*, SIAM Review **45** (2003), 167–256.

[14] E. Ravasz and A.L. Barabási, *Hierarchical organization in complex networks*, Physical Review E **67** (2003), 026112.

[15] D.J. Watts, *Six degrees: the science of a connected age*, W.W. Norton, 2003.

[16] G.U. Yule, *A mathematical theory of evolution, based on the conclusions of Dr. J. C. Willis*, Philos. Trans. Roy. Soc. London Ser. B **213** (1924), 21–87.

Department of Statistics, 367 Evans Hall, U.C. Berkeley, CA 94720
E-mail address: aldous@stat.berkeley.edu