

# An Internet Graph Model Based on Trade-Off Optimization

José Ignacio Alvarez-Hamelin<sup>1</sup> and Nicolas Schabanel<sup>2</sup>

<sup>1</sup> LRI, Université Paris Sud, Bât. 490, 91405 Orsay CEDEX, France, <http://www.lri.fr/~ihameli>.

<sup>2</sup> CNRS, LIP, ÉNS Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France, <http://perso.ens-lyon.fr/nicolas.schabanel/>.

Received: 11th June 2004/ Revised version: 11th June 2004

**Abstract** This paper presents a new model for the Internet graph (AS graph) based on the concept of heuristic trade-off optimization, introduced by Fabrikant, Koutsoupias and Papadimitriou in [5] to grow a random tree with a heavily tailed degree distribution. We propose here a generalization of this approach to generate a general graph, as a candidate for modeling the Internet. We present the results of our simulations and an analysis of the standard parameters measured in our model, compared with measurements from the physical Internet graph.

**PACS.** 89.75.-k Complex systems – 89.75.Hc Networks and genealogical trees – 89.75.Da Systems obeying scaling laws – 89.75.Fb Structures and organization in complex systems – 89.65.Gh Social and economic systems - Economics; econophysics, financial markets, business and management

## 1 Introduction

### 1.1 Motivations

The observations made by the three Faloutsos brothers [6] in 1999 were a striking revelation for the computer science community: Internet topology is not what was expected but follows “scale-free” properties, such as power-laws on the degree distribution and other parameters. These properties do not match the standard laws of the classic Erdős-Rényi random graph model. The models used to date for simulations, prediction and mathematical proofs do not match the “real networks”. “Scale-free” properties mean that the parameters of the system do not have typical values, in particular, very high values are reasonably probable (as opposed to Gaussian distributions). The distribution of the degrees (numbers of distinct connections) measured on the Internet autonomous systems (ASs)<sup>1</sup> graph follows a power law: the number of ASs connected to  $d$  other ASs is proportional to  $d^{-\beta}$  with  $\beta \approx 2.1$ . Interestingly enough, the value of  $\beta$  is stable over several years.

Ever since, much research has been conducted in order to try to explain these facts, by proposing various random graph models. These models are important to obtain better predictions or simulations of network behavior (such as virus spreading [11]), but also, hopefully, to build a new theoretical framework for more formal results.

<sup>1</sup> There are essentially two levels for studying the physical graph of the Internet: the router level (230.000 to 650.000 nodes, depending on the exploration) and the AS level (about 12 000 nodes). An autonomous system (AS) is an independent entity, such as a university, a local network... The AS graph is basically the hierarchical level along which packets are routed.

### 1.2 Previous Results

Different models aim to reproduce as closely as possible the power laws measured regularly on the Internet: examples are BRITE [10], INET [7],... These models are very sophisticated and follow reasonably the distributions observed (as long as one manages to compile their generators). However, these models do not offer satisfying explanations of the dynamics going on, and certainly do not allow the construction of a simple theoretical framework.

The first model leading to power law on degrees, known as linear preferential attachment, is due to Albert and Barabási [1]. In this model, nodes are inserted one after the other and every new node is connected to a fixed (or random) number of existing nodes, chosen with a probability proportional to their current degree. [1] shows that the distribution of the degrees  $d$  is a  $d^{-3}$  law. Another interesting possible explanation for the observed power laws was proposed by Kumar et al [9] for the Web graph (the directed graph of the HTML pages connected by hyperlinks). In their model, nodes are again inserted one after the other; each new node 1) is connected to a fixed (or random) number of existing nodes, chosen uniformly, and 2) copies some of the links of a fixed (or random) number of existing nodes, chosen again uniformly. This model yields also a power law on the degree distribution. More recently, inspired by the work of [3], Fabrikant, Koutsoupias and Papadimitriou in [5] propose a new model, which grows a random tree incrementally: every new node is connected to a pre-existing node that minimizes a trade-off between cost and efficiency. They show that the degree distribution of the nodes is heavy tailed, and raise the conjecture that power laws, often qualified “signature of human behavior”,

may be the product of greedy multi-criteria optimization. This question deserves to be studied in detail. The FKP tree model was further studied theoretically in [2] and [8].

### 1.3 Our Contribution

This article aims to evaluate how the optimization-based scheme suggested in [5] may effectively apply to the Internet. The first challenge is to obtain a model which gets closer to the dynamic observed for the Internet (see [4]) yet still remains simple. The second is to obtain a time-efficient generation. Our new model generalizes the model of [5] by growing a graph with a dynamic similar to the one observed by [4] for the Internet. We compare the result of an important simulation study using our model to the classical measures of Internet topology, and we propose an analysis of the parameters of our model to get as close as possible to these target values. We observe that parameters values in our model are close enough to those observed on the Internet, to validate the concept introduced in [5]. This gives a new plausible explanation for the power laws observed in the Internet structure. Our study ends with suggestions for improving our model, in particular by adding capacities to the links, which corresponds to a real request of the network community in order to conduct reliable simulations.

## 2 The Model

Our model is a generalization of the FKP random growing tree introduced by Fabrikant, Koutsoupias and Papadimitriou [5], that we will describe first.

### 2.1 The original FKP model

Inspired by the work of Carlson and Doyle on power laws and fault tolerance maximization (see [3]), the authors of [5] propose to build a random tree as follows. The sequence of nodes is a sequence of random points  $(x_i)$  chosen uniformly in the unit square  $[0, 1]^2$ . The first point  $x_0$  is the root of the tree, and every new point  $x_i$  gets connected to the pre-existing node  $x_j$ ,  $j < i$ , in the tree that minimizes:

$$d(x_i, x_j) + \theta \cdot h(x_j, x_0) \quad (1)$$

where  $\theta$  is the parameter of the model<sup>2</sup>, and where  $d(x_i, x_j)$  and  $h(x_j, x_0)$  are, respectively, the Euclidean distance between  $x_i$  and  $x_j$  in  $[0, 1]^2$ , i.e. the length of the link from  $x_i$  to  $x_j$  (the cost for drawing the link), and the hop distance from  $x_j$  to the root  $x_0$  in the tree, i.e. the “operating cost”. The tree grows by greedily optimizing a trade-off between two opposite costs: the cost for building the edge, and the operating cost of the resulting network (assuming that

<sup>2</sup> The original parameter in [5] is  $\alpha = 1/\theta$ . Adopting  $\theta$  improves readability, by first avoiding a lot of  $1/\alpha$  and, second, because  $\theta$  is the length parameter of the model.

nodes only communicate with or through the root). The model exhibits an interesting phase transition-like behavior for the degree distribution of the  $n$  first nodes of the tree: for  $\theta > 1/\sqrt{2}$ , one always get a star; for  $\theta < 1/4$ , as long as  $n \ll 1/\theta^2$ , the degree distribution is exponential; but as soon as  $n \gg 1/\theta^2$ , the degree distribution becomes heavy tailed. The authors raise the conjecture that “power law-like” distributions observed in human activities, may be the product of “balanced” trade-off optimization (referring to the fact that heavy tailed distribution appears for relatively “balanced” values of  $\theta$  relatively to  $n$ ). Kenyon and Schabanel [8] have shown the same kind of transition if one draws from every new node, two links, instead of one, to the two nodes  $x_j$  and  $x_{j'}$  that minimizes (1); let us denote this model FKP<sub>2</sub>. The FKP tree was studied further by Berger et al [2]. They show that its degree distribution did not follow a “standard power law” in the sense that it has way too many leaves ( $n - o(n)$ ). Furthermore, Kenyon and Schabanel [8] demonstrate that when  $\theta$  is constant, only a small number of nodes are grandfathers. The FKP model is thus, *as is*, a poor model for the Internet, but opens interesting questions: may trade-off optimization yield new and easy models for growing networks?

### 2.2 Our model

As seen above, the models FKP, which builds a tree, or FKP<sub>2</sub>, which has no degree 1 node, are not satisfying candidates to model the Internet. We extend the FKP approach to propose a new, more realistic model, following more closely the observed Internet dynamic. According to [4], AS graph dynamic from 1998 to 2000 had the following characteristics:

Fact 1: Most new nodes have degree 1 (87%) or 2 (12%) when created.

Fact 2: About 50% of the links appear independently from node creations.

Fact 3: One node (resp. a link) is deleted every third node creation (resp. second link creations) – which induces an exponential growth. Furthermore, AS (resp. link) deletions are strongly correlated with AS (resp. link) creations.

The correlation in Fact 3 suggests that most link or AS deletions are indeed network upgrades. We thus propose to neglect AS and link deletions in our model. As far as we know, none of the proposed models incorporates node or link deletions presently.

**Our model.** Our model constructs a graph accordingly to Facts 1 and 2. In all its generality, our model has five parameters  $(\theta, \gamma, k, q, \tau)$  that will be discussed below. The nodes are a sequence of  $n$  points  $(x_0, \dots, x_{n-1})$  randomly chosen in  $\mathbb{R}^2$  according to some probability distribution. We study in this article two distributions: the uniform distribution over  $[0, 1]^2$ , and the distribution of the US population given by the 1990 census. The construction

starts with the point  $x_0$ .  $x_0$  is the first origin of the graph. This origin, denoted  $\omega$ , changes from one node to another regularly over time in our model, in order to break the model centrality. The insertion of node  $x_i$ ,  $i \geq 1$ , takes three steps :

1. New node  $x_i$  is linked to the  $k$  distinct nodes  $x_{j_1}, \dots, x_{j_k}$  that minimize at the time of the insertion:

$$d(x_i, x_j) + \theta \cdot h(x_j, \omega) \quad (2)$$

2. Then,  $q$  new edges  $e_1, \dots, e_q$  are added one after the other. If  $G$  denotes the graph just after the insertion of edge  $e_{j-1}$ ,  $e_j$  is the edge that minimizes the following quantity:

$$\ell(e_j) + \frac{\gamma}{i} \sum_{a=0}^i (h_{G \cup e_j}(x_a, \omega) - h_G(x_a, \omega)) \quad (3)$$

where  $\ell(e_j)$  denotes the Euclidean length of  $e_j$  in the plane, and where  $h_G(x_a, \omega)$  and  $h_{G \cup e_j}(x_a, \omega)$  are the hop distances of node  $x_a$  to the origin  $\omega$  in graphs  $G$  and  $G \cup e_j$  respectively. The  $\gamma$ -weighted term in (3) is thus the expected decrease of the hop distance from the current nodes to the origin  $\omega$ .

3. Every  $\tau$  node insertions, change the origin to a random node  $x_a$ , chosen with a probability proportional to its degree:  $\omega := x_a$ .

Note that when  $k = 1$  (resp.  $k = 2$ ),  $q = 0$  and  $\tau = \infty$ , it is the FKP model (resp. FKP<sub>2</sub> model). While Step 1 determines how the new node joins the network, the origin improves at Step 2 its global connectivity to all the other nodes by adding a new edge in the network. Step 3, moves the origin randomly, which allows us to get an isotropic network for a very small amount of computation time – this basically approximates to a global optimization.

**Choosing  $k$ ,  $q$ ,  $\tau$  to model AS graph.** Fact 1 suggests that most of the newly created nodes have degree 1. We have thus chosen  $k = 1$  in our simulation (we could also have considered  $k$  as a random variable). Fact 2 suggests that half of the edges are created independently from node creation. These insertions take place in step 2 in our model. We have then chosen  $q = 1$ . Finally, in order to maximize our model isotropy, we chose  $\tau = 1$ : the origin changes after each iteration.

*Our model is left now with only two parameters  $\theta$  and  $\gamma$ .*

**Implementation.** Computational efficiency of our model is a central concern. Randomly changing the origin allows us to save a factor of  $n$  in the computation time: it allows us to obtain isotropy without having to minimize the hop-distance to all the other nodes in the trade-offs in steps 1 and 2. Step 2 is the most time consuming; our implementation takes presently  $O(n^4)$  time; this is a serious obstacle to generate relatively big graphs ( $> 2000$  nodes): generating a 1000 node graph takes about 15 min on a 1GHz

Powerbook G4 (1 hour for 1500 nodes). Considerable improvements in time performance can be achieved using hash tables by hop-distances to the origin, and Voronoï diagrams. The main goal of this paper is however to study the pertinence of this approach for Internet modeling, not yet its optimization. We are currently working on including these improvements to obtain  $O(n^{2+\epsilon})$  time generation of  $n$  node graphs. Figure 10 gives two examples of graphs generated by our model.

**The thermodynamic limit.** As noted in [5, 2] on the FKP model, unless  $\theta$  or  $\gamma$  depend on  $n$ , the network converges to a constant diameter graph with a highly connected “kernel” and a large majority of leaves. This article does not consider  $\theta$  and  $\gamma$  as functions of  $n$ ; the main reason is that our current implementation only allows us to reach relatively small values of  $n$  ( $n \lesssim 1500$ ) and thus adding a dependence on  $n$  would not make sense at this point. Again, this paper is an exploratory study of the validity of the trade-off optimization principle to model Internet growth, further improvements should be considered in the future.

### 3 Model validation

In order to validate our model we compare it with some widely used parameters. The reason we use these parameters is very empiric: most of the publications refer to them. We believe that pertinent parameters may differ from one application to another. It may even be the case that the best graph model (in the sense: the observed phenomenon behave the same as in reality, and is tractable on the model) for a given routing problem over the Internet may differ significantly from the actual Internet graph. Finding which parameters one should observe to compare two models given a particular application is presently an important open issue.

The parameters that we measured are:

- the average degree,  $\overline{deg}$  ( $4 - \epsilon$  trivially for our model).
- the clustering coefficient,  $\overline{Clust}$ : the average probability that two neighbors of a given node are directly connected to each other. (Note that this basically counts the number of triangles.)
- the minimum and average eccentricity,  $ex_m$  and  $\overline{ex}$  respectively:  $ex_m = \min_i \max_j h(x_i, x_j)$  and  $\overline{ex} = \max_j h(\cdot, x_j)$ .
- the graph diameter,  $D$ .
- the average hop distance,  $\overline{h(\cdot, \cdot)}$ .
- the exponent  $\beta$  of the closest power law to the degree distribution:  $\Pr\{\text{a node has degree } d\} \sim d^{-\beta}$ .
- the absolute correlation coefficient  $ACC$  of the power law on degrees. ( $|ACC| \in [0, 1]$ ; the correlation is perfect if  $|ACC| = 1$ , and inexistent if  $|ACC| = 0$ .)

### 3.1 Statistical properties of random graphs generated by our model

We first explore the possible values of the parameters  $\theta$  and  $\gamma$ , and then study the evolution of the measured parameters with the size of the graph. We observe that  $\theta$  and  $\gamma$  modify essentially the exponent  $\beta$  of the computed power law observed on the degree distribution:

$$\Pr\{\text{a node has degree } d\} \sim d^{-\beta} \quad (4)$$

**Influence of  $\gamma$ .** Our empirical measurements show that only  $\max(\gamma, \theta)$  seems to influence  $\beta$ : for instance, when  $\gamma < \theta$ , the graphs obtained are independent of  $\gamma$ . Therefore, we adopt  $\theta = \gamma$ . *We are now left with one single parameter,  $\theta$ , in our model.*

**Influence of  $\theta$  on the degree power law.**  $\theta$  is the length parameter of the model: it is basically the maximum length of a link one can draw. Naturally, the diameter and eccentricity decrease when  $\theta$  increases. The exponent  $\beta$  is also a decreasing function of  $\theta$ : as  $\theta$  increases, there are more and more high degree nodes. Surprisingly enough, the clustering coefficient also decreases when  $\theta$  increases: it is more profitable to connect closer to the current origin, than drawing shorter links towards its neighbors. As the graph gets a “smaller world” (smaller diameter), its clustering coefficient gets smaller too! This is an interesting peculiarity of these trade-off optimization based models. Figures 1–7 sum up these observations that will be discussed below.

**Influence of the size of the graph.** Table 1 shows the influence of the size  $n$  of the graph on the measured parameters, and compares them to the measured AS and IR graphs. We have generated graphs for two values of  $\theta$ :  $\theta = 0.03$  and  $\theta = 0.001$ . Both values yield a satisfying exponent  $\beta$  close to that the Internet (roughly  $\beta \simeq 2.1$  [12]).

We verified that as  $n$  increases, the exponent  $\beta$ , the clustering coefficient and the diameter get closer to AS graphs values. This may be due to the fact that the highest graph size we can reach is closer to the AS graph size than to the IR graph size. Hence, comparison with the IR may not be relevant.

Table 2 compares the parameters measured on FKP and FKP<sub>2</sub> models ( $q = 0$ ,  $\tau = \infty$ , and  $k = 1$  or  $2$  in our model) for the same value of  $\theta$ . We observe that the exponent  $\beta$  in the FKP model is closer to Internet’s  $\beta$ , but its clustering coefficient is null and the diameter is much larger, mainly because the graph is a tree. Concerning the FKP<sub>2</sub> model, we observe that there is no degree 1 node and that the clustering coefficient is too high compared to the AS graph. We conclude that our model fits better the real values observed on the Internet and improves considerably the previous approaches based on trade-off optimization.

**Statistical analysis.** In order to obtain a precise evolution of the parameters, we have undertaken a large simulation study. For every value of  $\theta = \gamma$  from 0.001 to 0.1 in steps of 0.001, we have generated 100 graphs of  $n = 1000$  nodes. For each value of  $\theta$ , we computed the average values and standard deviation of the different parameters. The results are shown in Figures 1–7. We observe that our model is stable and coherent, i.e., one can obtain any desired  $\beta$  value with an error of  $\pm 5\%$  approximately.

The length parameter  $\theta$  in the optimization process controls where a new node or a new link is connected to. When  $\theta$  (and  $\gamma$ ) is low, any new node connects to a node near by without paying too much attention to the distance to the current network origin. The distance to the network origin increases as the graph grows, and the diameter is larger. One observes that the diameter/eccentricity ratio remains almost constant. The node degree is rather low, because it depends only on the probability of having nodes near by.

When  $\theta$  (and  $\gamma$ ) is low, the situation for the new links is similar. They tend to connect nodes near to each other. Hence, the diameter does not significantly decrease. However, the clustering coefficient is high because the new links are created in the local environment, and the probability of building triangles is high.

On the other hand, when  $\theta$  (and  $\gamma$ ) is high, any new node tends to connect to a node close to the current network origin, in the sense of the hop distance, regardless of the Euclidean distance. The diameter of the graph is then lower. This also creates very popular nodes, which decreases  $\beta$ . Concerning the links, new extra links tend to connect popular nodes to distant nodes in order to decrease the average hop distance to the current origin. Therefore, triangles are less profitable, and the clustering coefficient decreases.

**Generation based on United States population.** Using the underlying geometry of our model, we have generated graphs where the nodes are distributed according the US population (based on the 1990 census). This approach is particularly relevant since Yook *et al.* [13] have shown correlation between the localization of the population and of the routers. We observe in Figures 8 and 9, and in Table 3, that the behavior is similar to the uniform distribution case on the unit square  $[0, 1]^2$  and gives satisfying pictures of the network.<sup>3</sup> Note that thanks to step 2 of our generation procedure, extra long links are created across the US (much longer than  $\theta$ ), in order to improve communication performance between the two coasts. This was not the case in the previous FKP and FKP<sub>2</sub> models. Note that the population distribution biases the degree distribution.

<sup>3</sup> See for instance:

<http://www.caida.org/tools/visualization/mapnet/>

$n$	$\theta = \gamma$	$\overline{deg}$	$\overline{Clust}$	$ex_m$	$\overline{ex}$	D	$h(\cdot, \cdot)$	$\beta$	$ACC$
20	0.001	3.50	0.2619	3	3.950	5	2.345	0.42	0.43
20	0.030	3.50	0.2444	3	3.850	5	2.235	0.72	0.75
50	0.001	3.80	0.2519	4	5.560	7	3.186	0.98	0.80
50	0.030	3.80	0.2305	4	5.400	7	3.064	0.88	0.76
100	0.001	3.90	0.2268	5	6.850	9	3.862	1.27	0.89
100	0.030	3.90	0.2217	5	6.660	9	3.663	1.27	0.87
200	0.001	3.95	0.2378	5	7.565	10	4.499	1.54	0.89
200	0.030	3.95	0.1992	5	6.830	9	4.129	1.45	0.95
500	0.001	3.98	0.2038	6	9.070	12	5.495	1.85	0.92
500	0.030	3.98	0.1910	5	7.480	9	4.864	1.77	0.94
1000	0.001	3.99	0.2186	7	10.436	14	6.240	1.94	0.93
1000	0.030	3.99	0.1887	6	8.510	11	5.377	2.02	0.94
2000	0.001	3.99	0.2186	7	10.436	14	6.240	2.14	0.93
2000	0.030	3.99	0.1798	6	8.773	11	5.821	2.12	0.96
AS graph		4.18	0.22	—	7	10	3.62	2.1	>0.96
IR graph		2.8	0.03	—	20	30	9.51	2.1 (?)	>0.96

**Table 1.** Measurements on our random graph models with  $k = 1$  and  $q = 1$ , compared with the measurements made on real AS and IR graphs.

$n$	$k$	$q$	$\overline{deg}$	$\overline{Clust}$	$ex_m$	$\overline{ex}$	D	$h(\cdot, \cdot)$	$\beta$	$ACC$
100	1	0	1.98	0	6	8.94	11	5.379	1.85	0.96
100	2	0	3.94	0.2563	4	5.61	7	3.714	1.92	0.95
200	1	0	1.99	0	6	9.36	12	5.884	1.93	0.97
200	2	0	3.97	0.2931	4	6.42	8	4.190	2.03	0.98
500	1	0	2	0	7	10.15	13	6.521	2.16	0.98
500	2	0	4	0.2495	5	7.218	9	4.886	2.21	0.98
1000	1	0	2	0	8	11.4	15	7.018	2.08	0.98
1000	2	0	4	0.2538	5	7.75	10	5.289	2.32	0.98

**Table 2.** Measurements on random FKP graphs ( $k = 1$ ), and FKP<sub>2</sub> graphs ( $k = 2$ ) with  $\theta = 0.03$ .

$n$	$k$	$q$	$\overline{deg}$	$\overline{Clust}$	$ex_m$	$\overline{ex}$	D	$h(\cdot, \cdot)$	$\beta$	$ACC$
1000	1	1	3.99	0.2081	5	7.103	9	4.364	2.46	0.91

**Table 3.** Measurements on our random graph model based on the US population distribution, with  $\theta = \gamma = 0.03$ .

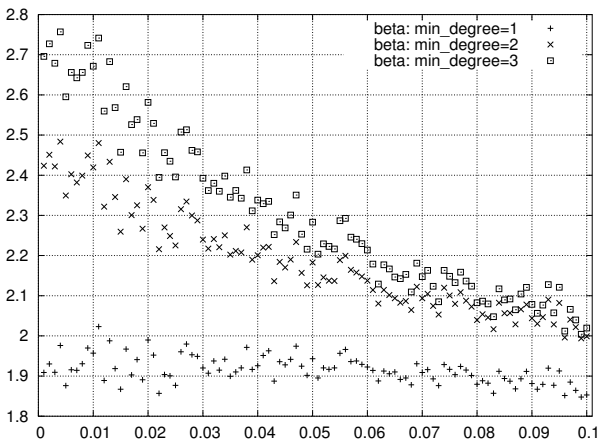
## 4 Conclusion and Summary

In this article, we have shown the relevance of a multi-criteria greedy optimization approach as suggested in [5], to model Internet topology. This multicriteria greedy optimization approach provides a plausible economical explanation of the observations made of the Internet, such as the power law on degrees. Another interesting point is the natural geographic representation of this model which makes it easy to read and interpret. Note also that it is still sufficiently simple, so that theoretical descriptions may not be out of reach. Some peculiarities of our model, such as the clustering coefficient increasing as the diameter grows, may also help to determine the relationships between the parameters: small world does not necessarily imply high clustering coefficient.

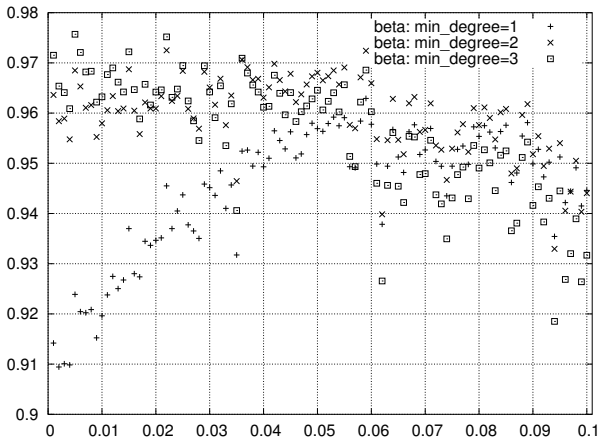
Currently, the main weakness of our model is its generation time. We are currently working on a new implementation taking advantage of the underlying geometry

to decrease the processing time from  $O(n^4)$  time (present implementation) down to  $O(n^{2+\epsilon})$ . Another issue is the relatively small number of degree 1 nodes in comparison with the Internet.


We also want to stress that the current parameters used to validate network models may not be appropriate. Some relationships between them are known. But the individual relevance of the parameter probably depends on the target application. It is also clear that because of its technological complexity and dynamic, the Internet will certainly not be modeled by one concept alone (preferential attachment [1], imitation [9], or trade-off optimization [5],...). For instance, these different processes may influence the structure of the Internet at different levels: e.g., trade-off optimization at the router level, preferential attachment at the AS level, and imitation at the web level; we may also have to consider other kind of levels. One will probably need to mix these concepts with others in order to obtain a satisfying behavior.



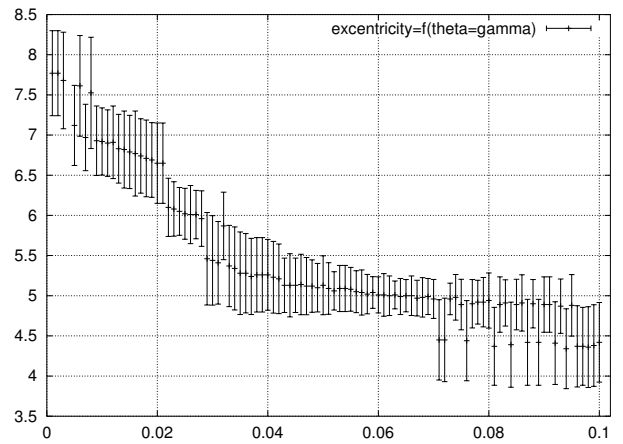
**Figure 1.** The exponent  $\beta$  as a function of  $\theta$  in our model.  $\beta$  is computed in three ways: 1) on all the nodes, 2) on nodes with degree  $> 1$ , and 3) on nodes with degree  $> 2$ . The largest standard deviations are  $\pm 10\%$ , and typical values are  $\pm 5\%$ .




**Figure 2.** The ACC for exponent  $\beta$  as a function of  $\theta$  in our model. The largest standard deviations are  $\pm 2\%$ , and typical values are  $\pm 1\%$ .



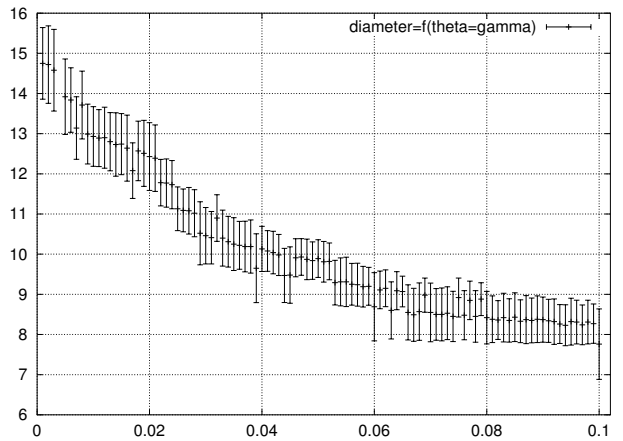
**Figure 3.** Clustering coefficient as a function of  $\theta$  in our model.



**Figure 4.** Eccentricity as a function of  $\theta$  in our model.



**Figure 5.** Average eccentricity as a function of  $\theta$  in our model.




**Figure 6.** Diameter as a function of  $\theta$  in our model.


We would like to thank Claire Kenyon for very useful and interesting discussions.

## References


1. R. Albert and A.-L. Barabasi. Topology of evolving networks: Local events and universality. *Physical Review Letters*, 85:5234–5237, 2000.
2. N. Berger, Béla Bollobás, C. Borgs, J. Chayes, and O. Riordan. Degree distribution of the FKP network model. In *Proc. of ICALP 2003*, volume 2719, pages 725–738, 2003.
3. J. M. Carlson and J. Doyle. Highly optimized tolerance: a mechanism for power laws in designed systems. *Physics Review E*, 60(2):1412–1427, 1999.
4. Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited. In *IEEE INFOCOM 2002*, pages 608–617, New York, NY, June 2002.
5. Alex Fabrikant, Elias Koutsoupias, and Christos H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the internet. *LNCS*, 2380:110–, Jun 2002.
6. Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262. ACM Press, 1999.
7. C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR443-00, Department of EECS, University of Michigan, 2000.
8. C. Kenyon and N. Schabanel. Finer analysis of the FKP random graph model for the Internet. Personal communication, 2003.
9. R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the Web graph. In *Proc. 41st Ann. Symp. on Foundation of Computer Science (FOCS)*, pages 57–65. IEEE Computer Society, 2000.
10. A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *Computer Communications Review*, 30(2):18–28, Apr 2000.
11. Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86:3200–3203, 2001.
12. Alexei Vazquez, Romualdo Pastor-Satorras, and Alessandro Vespignani. Internet topology at the router and autonomous system level. *Physical Review Letters*, 87:258701–, 2001.
13. S. Yook, H. Jeong, and A. L. Barabasi. Modeling the internet’s large-scale topology. *Proc. Nat. Acad. Sci.*, 99:13382–13386, 2002.




**Figure 7.** Average hop distance as a function of  $\theta$  in our model.




**Figure 8.** An example of a random graph generated by our model with the US population distribution, with  $n = 1000$  nodes,  $k = 1$ ,  $q = 1$ , and  $\theta = \gamma = 0.03$ . Each node is represented by a small disc whose diameter is proportional to its degree.



**Figure 9.** Observed power law on degrees in our random graph models based on US population distribution, with  $n = 1000$  nodes,  $k = 1$ ,  $q = 1$ , and  $\theta = \gamma = 0.03$ .



$$\theta = \gamma = 0.001.$$



$$\theta = \gamma = 0.03.$$

**Figure 10.** Two generated graphs of 500 nodes, with  $k = 1$  and  $q = 1$  and different values for  $\theta$  and  $\gamma$ .