

Realistic, Mathematically Tractable Graph Generation and Evolution, Using Kronecker Multiplication

Jurij Leskovec¹, Deepayan Chakrabarti¹, Jon Kleinberg², and
Christos Faloutsos¹

¹ School of Computer Science, Carnegie Mellon University
{jure, deepay, christos}@cs.cmu.edu

² Department of Computer Science, Cornell University
kleinber@cs.cornell.edu

Abstract. How can we generate realistic graphs? In addition, how can we do so with a mathematically tractable model that makes it feasible to analyze their properties rigorously? Real graphs obey a long list of surprising properties: Heavy tails for the in- and out-degree distribution; heavy tails for the eigenvalues and eigenvectors; small diameters; and the recently discovered “Densification Power Law” (DPL). All published graph generators either fail to match several of the above properties, are very complicated to analyze mathematically, or both. Here we propose a graph generator that is mathematically tractable and matches this collection of properties. The main idea is to use a non-standard matrix operation, the *Kronecker product*, to generate graphs that we refer to as “Kronecker graphs”.

We show that Kronecker graphs naturally obey all the above properties; in fact, we can rigorously *prove* that they do so. We also provide empirical evidence showing that they can mimic very well several real graphs.

1 Introduction

What do real graphs look like? How do they evolve over time? How can we generate synthetic, but realistic, time-evolving graphs? Graph mining has been attracting much interest recently, with an emphasis on finding patterns and abnormalities in social networks, computer networks, e-mail interactions, gene regulatory networks, and many more. Most of the work focuses on static snapshots of graphs, where fascinating “laws” have been discovered, including small diameters and heavy-tailed degree distributions.

Work partially supported by the National Science Foundation under Grants No. IIS-0209107, SENSOR-0329549, IIS-0326322, CNS-0433540, CCF-0325453, IIS-0329064, CNS-0403340, CCR-0122581, a David and Lucile Packard Foundation Fellowship, and also by the Pennsylvania Infrastructure Technology Alliance (PITA), a partnership of Carnegie Mellon, Lehigh University and the Commonwealth of Pennsylvania’s Department of Community and Economic Development (DCED). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, or other funding parties.

A realistic graph generator is important for at least two reasons. The first is that it can generate graphs for extrapolations, “what-if” scenarios, and simulations, when real graphs are difficult or impossible to collect. For example, how well will a given protocol run on the Internet five years from now? Accurate graph generators can produce more realistic models for the future Internet, on which simulations can be run. The second reason is more subtle: it forces us to think about the patterns that a graph generator should obey, to be realistic.

The main contributions of this paper are the following:

- We provide a generator which obeys all the main static patterns that have appeared in the literature.
- Generator also obeys the recently discovered temporal evolution patterns.
- Contrary to other generators that match this combination of properties, our generator leads to tractable analysis and rigorous proofs.

Our generator is based on a non-standard matrix operation, the *Kronecker product*. There are several theorems on Kronecker products, which actually correspond exactly to a significant portion of what we want to prove: heavy-tailed distributions for in-degree, out-degree, eigenvalues, and eigenvectors. We also demonstrate how a Kronecker Graph can match the behavior of several real graphs (patent citations, paper citations, and others). While Kronecker products have been studied by the algebraic combinatorics community (see e.g. [10]), the present work is the first to employ this operation in the design of network models to match real datasets.

The rest of the paper is organized as follows: Section 2 surveys the related literature. Section 3 gives the proposed method. We present the experimental results in Section 4, and we close with some discussion and conclusions.

2 Related Work

First, we will discuss the commonly found (static) patterns in graphs, then some recent patterns on temporal evolution, and finally, the state of the art in graph generation methods.

Static Graph Patterns: While many patterns have been discovered, two of the principal ones are heavy-tailed degree distributions and small diameters.

Degree distribution: The degree-distribution of a graph is a power law if the number of nodes c_k with degree k is given by $c_k \propto k^{-\gamma}$ ($\gamma > 0$) where γ is called the power-law exponent. Power laws have been found in the Internet [13], the Web [15, 7], citation graphs [24], online social networks [9] and many others. Deviations from the power-law pattern have been noticed [23], which can be explained by the “DGX” distribution [5]. DGX is closely related to a truncated lognormal distribution.

Small diameter: Most real-world graphs exhibit relatively small diameter (the “small-world” phenomenon): A graph has diameter d if every pair of nodes can be connected by a path of length at most d . The diameter d is susceptible to outliers. Thus, a more robust measure of the pairwise distances between nodes of a graph is the *effective diameter* [26]. This is defined as the minimum number

of hops in which some fraction (or quantile q , say $q = 90\%$) of all connected pairs of nodes can reach each other. The effective diameter has been found to be small for large real-world graphs, like Internet, Web, and social networks [2, 21].

Scree plot: This is a plot of the eigenvalues (or singular values) of the adjacency matrix of the graph, versus their rank, using a log-log scale. The scree plot is also often found to approximately obey a power law. The distribution of eigenvector components (indicators of “network value”) has also been found to be skewed [9].

Apart from these, several other patterns have been found, including the “stress” [14, 9], “resilience” [2, 22], “clustering coefficient” and many more.

Temporal evolution Laws: Densification and shrinking diameter: Two very recent discoveries, both regarding time-evolving graphs, are worth mentioning [18]: (a) the “effective diameter” of graphs tends to shrink or stabilize as the graph grows with time, and (b) the number of edges $E(t)$ and nodes $N(t)$ seems to obey the *densification power law* (DPL), which states that

$$E(t) \propto N(t)^a \quad (1)$$

The *densification exponent* a is typically greater than 1, implying that the average degree of a node in the graph is *increasing* over time. This means that real graphs tend to sprout many more edges than nodes, and thus are densifying as they grow.

Graph Generators: The earliest probabilistic generative model for graphs was a random graph model, where each pair of nodes has an identical, independent probability of being joined by an edge [11]. The study of this model has led to a rich mathematical theory; however, this generator produces graphs that fail to match real-world networks in a number of respects (for example, it does not produce heavy-tailed degree distributions).

The vast majority of recent models involve some form of *preferential attachment* [1, 2, 28, 15, 16]: new nodes join the graph at each time step, and preferentially connect to existing nodes with high degree (the “rich get richer”). This simple behavior leads to power-law tails and to low diameters. The diameter in this model grows slowly with the number of nodes N , which violates the “shrinking diameter” property mentioned above.

Another family of graph-generation methods strives for small diameter, like the *small-world* generator [27] and the Waxman generator [6]. A third family of methods show that heavy tails emerge if nodes try to optimize their connectivity under resource constraints [8, 12].

Summary: Most current generators focus on only one (static) pattern, and neglect the others. In addition, it is usually hard to prove properties of them. The generator we describe in the next section addresses these issues.

3 Proposed Method

The method we propose is based on a recursive construction. Defining the recursion properly is somewhat subtle, as a number of standard, related graph con-

struction methods fail to produce graphs that densify according to the patterns observed in practice, and they also produce graphs whose diameters increase. To produce densifying graphs with constant diameter, and thereby match the qualitative behavior of real network datasets, we develop a procedure that is best described in terms of the *Kronecker product* of matrices. To help in the description of the method, the accompanying table provides a list of symbols and their definitions.

Symbol	Definition
G_1	the initiator of a Kronecker Graph
N_1	number of nodes in initiator
E_1	number of edges in initiator
$G_1^{[k]} = G_k$	the k^{th} Kronecker power of G_1
a	densification exponent
d	diameter of a graph
$\mathcal{P}_1$	probability matrix

3.1 Main idea

The main idea is to create self-similar graphs, recursively. We begin with an *initiator* graph G_1 , with N_1 nodes and E_1 edges, and by recursion we produce successively larger graphs $G_2 \dots G_n$ such that the k^{th} graph G_k is on $N_k = N_1^k$ nodes. If we want these graphs to exhibit a version of the Densification Power Law, then G_k should have $E_k = E_1^k$ edges. This is a property that requires some care in order to get right, as standard recursive constructions (for example, the traditional Cartesian product or the construction of [4]) do not satisfy it.

It turns out that the *Kronecker product* of two matrices is the perfect tool for this goal. The Kronecker product is defined as follows:

Definition 1 (Kronecker product of matrices). *Given two matrices $\mathbf{A} = [a_{i,j}]$ and $\mathbf{B}$ of sizes $n \times m$ and $n' \times m'$ respectively, the Kronecker product matrix $\mathbf{C}$ of dimensions $(n * n') \times (m * m')$ is given by*

$$\mathbf{C} = \mathbf{A} \otimes \mathbf{B} \doteq \begin{pmatrix} a_{1,1}\mathbf{B} & a_{1,2}\mathbf{B} & \dots & a_{1,m}\mathbf{B} \\ a_{2,1}\mathbf{B} & a_{2,2}\mathbf{B} & \dots & a_{2,m}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n,1}\mathbf{B} & a_{n,2}\mathbf{B} & \dots & a_{n,m}\mathbf{B} \end{pmatrix} \quad (2)$$

We define the Kronecker product of two graphs as the Kronecker product of their adjacency matrices.

Observation 1 (Edges in Kronecker-multiplied graphs)

$$\text{Edge } (X_{ij}, X_{kl}) \in G \otimes H \text{ iff } (X_i, X_k) \in G \text{ and } (X_j, X_l) \in H$$

where X_{ij} and X_{kl} are nodes in $G \otimes H$, and X_i, X_j, X_k and X_l are the corresponding nodes in G and H , as in Figure 1.




Fig. 1. *Example of Kronecker multiplication:* Top: a “3-chain” and its Kronecker product with itself; each of the X_i nodes gets expanded into 3 nodes, which are then linked using Observation 1. Bottom row: the corresponding adjacency matrices, along with matrix for the fourth Kronecker power G_4 .

The last observation is subtle, but crucial, and deserves elaboration: Figure 1(a–c) shows the recursive construction of $G \otimes H$, when $G = H$ is a 3-node path. Consider node $X_{1,2}$ in Figure 1(c): It belongs to the H graph that replaced node X_1 (see Figure 1(b)), and in fact is the X_2 node (i.e., the center) within this small H -graph.

We propose to produce a growing sequence of graphs by iterating the Kronecker product:

Definition 2 (Kronecker power). *The k^{th} power of G_1 is defined as the matrix $G_1^{[k]}$ (abbreviated to G_k), such that:*

$$G_1^{[k]} = G_k = \underbrace{G_1 \otimes G_1 \otimes \dots \otimes G_1}_{k \text{ times}} = G_{k-1} \otimes G_1$$

The self-similar nature of the Kronecker graph product is clear: To produce G_k from G_{k-1} , we “expand” (replace) each node of G_{k-1} by converting it into a copy of G , and we join these copies together according to the adjacencies in G_{k-1} (see Figure 1). This process is very natural: one can imagine it as positing that communities with the graph grow recursively, with nodes in the community recursively getting expanded into miniature copies of the community. Nodes in the subcommunity then link among themselves and also to nodes from different communities.

3.2 Theorems and Proofs

We shall now discuss the properties of Kronecker graphs, specifically, their degree distributions, diameters, eigenvalues, eigenvectors, and time-evolution. Our ability to prove analytical results about all of these properties is a major advantage of Kronecker graphs over other generators. The next few theorems prove that several distributions of interest are multinomial for our Kronecker graph model. This is important, because a careful choice of the initial graph G_1 can make the resulting multinomial distribution to behave like a power-law or DGX distribution.

Theorem 1 (Multinomial degree distribution). *Kronecker graphs have multinomial degree distributions, for both in- and out-degrees.*

Proof. Let the initiator G_1 have the degree sequence $d_1, d_2, \dots, d_{N_1}$. Kronecker multiplication of a node with degree d expands it into N_1 nodes, with the corresponding degrees being $d \times d_1, d \times d_2, \dots, d \times d_{N_1}$. After Kronecker powering, the degree of each node in graph G_k is of the form $d_{i_1} \times d_{i_2} \times \dots \times d_{i_k}$, with $i_1, i_2, \dots, i_k \in (1 \dots N_1)$, and there is one node for each ordered combination. This gives us the multinomial distribution on the degrees of G_k . Note also that the degrees of nodes in G_k can be expressed as the k^{th} Kronecker power of the vector $(d_1, d_2, \dots, d_{N_1})$. **QED**

Theorem 2 (Multinomial eigenvalue distribution). *The Kronecker graph G_k has a multinomial distribution for its eigenvalues.*

Proof. Let G_1 have the eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_{N_1}$. By properties of the Kronecker multiplication [19, 17], the eigenvalues of G_k are k^{th} Kronecker power of the vector $(\lambda_1, \lambda_2, \dots, \lambda_{N_1})$. As in Theorem 1, the eigenvalue distribution is a multinomial. **QED**

A similar argument using properties of Kronecker matrix multiplication shows the following.

Theorem 3 (Multinomial eigenvector distribution). *The components of each eigenvector of the Kronecker graph G_k follow a multinomial distribution.*

We have just covered several of the static graph patterns. Notice that the proofs were direct consequences of the Kronecker multiplication properties.

Next we continue with the temporal patterns: the densification power law, and shrinking/stabilizing diameter.

Theorem 4 (DPL). *Kronecker graphs follow the Densification Power Law (DPL) with densification exponent $a = \log(E_1)/\log(N_1)$.*

Proof. Since the k^{th} Kronecker power G_k has $N_k = N_1^k$ nodes and $E_k = E_1^k$ edges, it satisfies $E_k = N_k^a$, where $a = \log(E_1)/\log(N_1)$. The crucial point is that this exponent a is independent of k , and hence the sequence of Kronecker powers follows an exact version of the Densification Power Law. **QED**

We now show how the Kronecker product also preserves the property of constant diameter, a crucial ingredient for matching the diameter properties of many real-world network datasets. In order to establish this, we will assume that the initiator graph G_1 has a self-loop on every node; otherwise, its Kronecker powers may in fact be disconnected.

Lemma 1. *If G and H each have diameter at most d , and each has a self-loop on every node, then the Kronecker product $G \otimes H$ also has diameter at most d .*

Proof. Each node in $G \otimes H$ can be represented as an ordered pair (v, w) , with v a node of G and w a node of H , and with an edge joining (v, w) and (x, y) precisely when (v, x) is an edge of G and (w, y) is an edge of H . Now, for an arbitrary pair of nodes (v, w) and (v', w') , we must show that there is a path of length at most d connecting them. Since G has diameter at most d , there is a path $v = v_1, v_2, \dots, v_r = v'$, where $r \leq d$. If $r < d$, we can convert this into a path $v = v_1, v_2, \dots, v_d = v'$ of length exactly d , by simply repeating v' at the end for $d - r$ times. By an analogous argument, we have a path $w = w_1, w_2, \dots, w_d = w'$. Now by the definition of the Kronecker product, there is an edge joining (v_i, w_i) and (v_{i+1}, w_{i+1}) for all $1 \leq i \leq d - 1$, and so $(v, w) = (v_1, w_1), (v_2, w_2), \dots, (v_d, w_d) = (v', w')$ is a path of length d connecting (v, w) to (v', w') , as required. **QED**

Theorem 5. *If G_1 has diameter d and a self-loop on every node, then for every k , the graph G_k also has diameter d .*

Proof. This follows directly from the previous lemma, combined with induction on k . **QED**

We also consider the *effective diameter* d_e ; we define the q -effective diameter as the minimum d_e such that, for at least a q fraction of the reachable node pairs, the path length is at most d_e . The q -effective diameter is a more robust quantity than the diameter, the latter being prone to the effects of degenerate structures in the graph (e.g. very long chains); however, the q -effective diameter and diameter tend to exhibit qualitatively similar behavior. For reporting results in subsequent sections, we will generally consider the q -effective diameter with $q = .9$, and refer to this simply as the *effective diameter*.

Theorem 6 (Effective Diameter). *If G_1 has diameter d and a self-loop on every node, then for every q , the q -effective diameter of G_k converges to d (from below) as k increases.*

Proof. To prove this, it is sufficient to show that for two randomly selected nodes of G_k , the probability that their distance is d converges to 1 as k goes to infinity.

We establish this as follows. Each node in G_k can be represented as an ordered sequence of k nodes from G_1 , and we can view the random selection of a node in G_k as a sequence of k independent random node selections from G_1 . Suppose that $v = (v_1, \dots, v_k)$ and $w = (w_1, \dots, w_k)$ are two such randomly selected nodes from G_k . Now, if x and y are two nodes in G_1 at distance d (such a pair (x, y)

exists since G_1 has diameter d), then with probability $1 - (1 - 2/N_1)^k$, there is some index j for which $\{v_j, w_j\} = \{x, y\}$. If there is such an index, then the distance between v and w is d . As the expression $1 - (1 - 2/N_1)^k$ converges to 1 as k increases, it follows that the q -effective diameter is converging to d . **QED**

3.3 Stochastic Kronecker Graphs

While the Kronecker power construction discussed thus far yields graphs with a range of desired properties, its discrete nature produces “staircase effects” in the degrees and spectral quantities, simply because individual values have large multiplicities. Here we propose a stochastic version of Kronecker graphs that eliminates this effect. counterparts.

We start with an $N_1 \times N_1$ *probability matrix* $\mathcal{P}_1$: the value p_{ij} denotes the probability that edge (i, j) is present. We compute its k^{th} Kronecker power $\mathcal{P}_1^{[k]} = \mathcal{P}_k$; and then for each entry p_{uv} of $\mathcal{P}_k$, we include an edge between nodes u and v with probability $p_{u,v}$. The resulting binary random matrix $R = R(\mathcal{P}_k)$ will be called the *instance matrix* (or *realization matrix*).

In principle one could try choosing each of the N_1^2 parameters for the matrix $\mathcal{P}_1$ separately. However, we reduce the number of parameters to just two: α and β . Let G_1 be the initiator matrix (binary, deterministic); we create the corresponding probability matrix $\mathcal{P}_1$ by replacing each “1” and “0” of G_1 with α and β respectively ($\beta \leq \alpha$). The resulting probability matrices maintain — with some random noise — the self-similar structure of the Kronecker graphs in the previous subsection (which, for clarity, we call *deterministic Kronecker graphs*).

We find empirically that the random graphs produced by this model continue to exhibit the desired properties of real datasets, and without the staircase effect of the deterministic version. The task of setting α and β to match observed data is a very promising research direction, outside the scope of this paper. In our experiments in the upcoming sections, we use heuristics which we describe there.

4 Experiments

Now, we demonstrate the ability of Kronecker graphs to match the patterns of real-world graphs. The datasets we use are:

- *arXiv*: This is a citation graph for high-energy physics research papers, with a total of $N = 29,555$ papers and $E = 352,807$ citations. We follow its evolution from January 1993 to April 2003, with one data-point per month.
- *Patents*: This is a U.S. patent citation dataset that spans 37 years from January 1963 to December 1999. The graph contains a total of $N = 3,942,825$ patents and $E = 16,518,948$ citations. Citation graphs are normally considered as directed graphs. For the purpose of this work we think of them as undirected.
- *Autonomous systems*: We also analyze a static dataset consisting of a single snapshot of connectivity among Internet autonomous systems from January 2000, with $N = 6,474$ and $E = 26,467$.

We observe two kinds of graph patterns — “static” and “temporal.” As mentioned earlier, common static patterns include the degree distribution, the scree plot (eigenvalues of graph adjacency matrix vs. rank), principal eigenvector of adjacency matrix and the distribution of connected components. Temporal patterns include the diameter over time, the size of the giant component over time, and the densification power law. For the diameter computation, we use a smoothed version of the effective diameter that is qualitatively similar to the standard effective diameter, but uses linear interpolation so as to take on non-integer values; see [18] for further details on this calculation.

Results are shown in Figures 2 and 3 for the graphs which evolve over time (*arXiv* and *Patents*). For brevity, we show the plots for only two static and two temporal patterns. We see that the deterministic Kronecker model already captures the qualitative structure of the degree and eigenvalue distributions, as well as the temporal patterns represented by the Densification Power Law and the stabilizing diameter. However, the deterministic nature of this model results in “staircase” behavior, as shown in scree plot for the deterministic Kronecker graph of Figure 2 (second row, second column). We see that the Stochastic Kronecker Graphs smooth out these distributions, further matching the qualitative structure of the real data; they also match the shrinking-before-stabilization trend of the diameters of real graphs.

For the Stochastic Kronecker Graphs we need to estimate the parameters α and β defined in the previous section. This leads to interesting questions whose full resolution lies beyond the scope of the present paper; currently, we searched by brute force over (the relatively small number of) possible initiator graphs of up to five nodes, and we then chose α and β so as to match well the edge density, the maximum degree, the spectral properties, and the DPL exponent.

Finally, Figure 4 shows plots for the static patterns in the *Autonomous systems* graph. Recall that we analyze a single, static snapshot in this case. In addition to the degree distribution and scree plot, we also show two typical plots [9]: the distribution of *network values* (principal eigenvector components, sorted, versus rank) and the *hop-plot* (the number of reachable pairs $P(h)$ within h hops or less, as a function of the number of hops h).

5 Observations and Conclusions

Here we list several of the desirable properties of the proposed Kronecker Graphs and Stochastic Kronecker Graphs.

Generality: Stochastic Kronecker Graphs include several other generators, as special cases: For $\alpha=\beta$, we obtain an Erdős-Rényi random graph; for $\alpha=1$ and $\beta=0$, we obtain a deterministic Kronecker graph; setting the G_1 matrix to a 2x2 matrix, we obtain the RMAT generator [9]. In contrast to Kronecker graphs, the RMAT cannot extrapolate into the future, since it needs to know the number of edges to insert. Thus, it is incapable of obeying the “densification law”.

Phase transition phenomena: The Erdős-Rényi graphs exhibit phase transitions [11]. Several researchers argue that real systems are “at the edge of




Fig. 2. *arXiv dataset*: Patterns from the real graph (top row), the deterministic Kronecker graph with G_1 being a star graph with 3 satellites (middle row), and the Stochastic Kronecker graph ($\alpha = 0.41$, $\beta = 0.11$ – bottom row). *Static* patterns: (a) is the PDF of degrees in the graph (log-log scale), and (b) the distribution of eigenvalues (log-log scale). *Temporal* patterns: (c) gives the effective diameter over time (linear-linear scale), and (d) is the number of edges versus number of nodes over time (log-log scale). Notice that the Stochastic Kronecker Graph qualitatively matches all the patterns very well.

chaos” [3, 25]. It turns out that Stochastic Kronecker Graphs also exhibit phase transitions. For small values of α and β , Stochastic Kronecker Graphs have many small disconnected components; for large values they have a giant component with small diameter. In between, they exhibit behavior suggestive of a phase transition: For a carefully chosen set of (α, β) , the diameter is large, and a giant component just starts emerging. We omit the details, for lack of space.

Theory of random graphs: All our theorems are for the deterministic Kronecker Graphs. However, there is a lot of work on the properties of random matrices (see e.g. [20]), which one could potentially apply in order to prove properties of the Stochastic Kronecker Graphs.

In conclusion, the main contribution of this work is a family of graph generators, using a non-traditional matrix operation, the *Kronecker product*. The resulting graphs (a) have all the static properties (heavy-tailed degree distribution, small diameter), (b) all the temporal properties (densification, shrinking diameter), and in addition, (c) we can formally prove all of these properties.

Several of the proofs are extremely simple, thanks to the rich theory of Kronecker multiplication. We also provide proofs about the diameter and “effective diameter”, and we show that Stochastic Kronecker Graphs can be tuned to mimic real graphs well.




Fig. 3. *Patents*: Again, Kronecker graphs match all of these patterns. We show only the Stochastic Kronecker graph for brevity.

References

1. R. Albert and A.-L. Barabasi. Emergence of scaling in random networks. *Science*, pages 509–512, 1999.
2. R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 2002.
3. P. Bak. How nature works : The science of self-organized criticality, Sept. 1996.
4. A.-L. Barabasi, E. Ravasz, and T. Vicsek. Deterministic scale-free networks. *Physica A*, 299:559–564, 2001.
5. Z. Bi, C. Faloutsos, and F. Korn. The DGX distribution for mining massive, skewed data. In *KDD*, pages 17–26, 2001.
6. B.M. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, 6(9), December 1988.
7. A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graph structure in the web: experiments and models. In *Proceedings of World Wide Web Conference*, 2000.
8. 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.
9. D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-MAT: A recursive model for graph mining. In *SDM*, 2004.
10. T. Chow. The Q-spectrum and spanning trees of tensor products of bipartite graphs. *Proc. Amer. Math. Soc.*, 125:3155–3161, 1997.
11. P. Erdos and A. Renyi. On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Science*, 5:17–67, 1960.
12. A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the internet (extended abstract), 2002.
13. M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.
14. M. Girvan and M. E. J. Newman. Community structure in social and biological networks. In *PNAS*, volume 99, 2002.




Fig. 4. *Autonomous systems:* Real (top) versus Kronecker (bottom). Columns (a) and (b) show the degree distribution and the scree plot, as before. Columns (c) and (d) show two more static patterns (see text). Notice that, again, the Stochastic Kronecker Graph matches well the properties of the real graph.

15. J. M. Kleinberg, S. R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. The web as a graph: Measurements, models and methods. In *Proceedings of the International Conference on Combinatorics and Computing*, 1999.
16. S. R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. Extracting large-scale knowledge bases from the web. *VLDB*, pages 639–650, 1999.
17. A. N. Langville and W. J. Stewart. The Kronecker product and stochastic automata networks. *Journal of Computation and Applied Mathematics*, 167:429–447, 2004.
18. J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *KDD '05*, Chicago, IL, USA, 2005.
19. C. F. V. Loan. The ubiquitous Kronecker product. *Journal of Computation and Applied Mathematics*, 123:85–100, 2000.
20. M. Mehta. *Random Matrices*. Academic Press, 2nd edition, 1991.
21. S. Milgram. The small-world problem. *Psychology Today*, 2:60–67, 1967.
22. C. R. Palmer, P. B. Gibbons, and C. Faloutsos. Anf: A fast and scalable tool for data mining in massive graphs. In *SIGKDD*, Edmonton, AB, Canada, 2002.
23. D. M. Pennock, G. W. Flake, S. Lawrence, E. J. Glover, and C. L. Giles. Winners don't take all: Characterizing the competition for links on the web. *Proceedings of the National Academy of Sciences*, 99(8):5207–5211, 2002.
24. S. Redner. How popular is your paper? an empirical study of the citation distribution. *European Physical Journal B*, 4:131–134, 1998.
25. R. Sole and B. Goodwin. *Signs of Life: How Complexity Pervades Biology*. Perseus Books Group, New York, NY, 2000.
26. S. L. Tauro, C. Palmer, G. Siganos, and M. Faloutsos. A simple conceptual model for the internet topology. In *Global Internet, San Antonio, Texas*, 2001.
27. D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.
28. J. Winick and S. Jamin. Inet-3.0: Internet Topology Generator. Technical Report CSE-TR-456-02, University of Michigan, Ann Arbor, 2002.