

- [1] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47-97, Jan. 2002. [[bib](#) | [.html](#)]

Complex networks describe a wide range of systems in nature and society, much quoted examples including the cell, a network of chemicals linked by chemical reactions, or the Internet, a network of routers and computers connected by physical links. While traditionally these systems were modeled as random graphs, it is increasingly recognized that the topology and evolution of real networks is governed by robust organizing principles. Here we review the recent advances in the field of complex networks, focusing on the statistical mechanics of network topology and dynamics. After reviewing the empirical data that motivated the recent interest in networks, we discuss the main models and analytical tools, covering random graphs, small-world and scale-free networks, as well as the interplay between topology and the network's robustness against failures and attacks.

- [2] R. Albert, H. Jeong, and A.-L. Barabási. Error and attack tolerance in complex networks. *Nature*, 406:387-482, 2000. [[bib](#) | [.html](#)]
- [3] D. Alderson, J. Doyle, R. Govindan, and W. Willinger. Toward an optimization-driven framework for designing and generating realistic internet topologies. In *ACM HotNets-I*, Princeton, NJ, Oct. 2002. [[bib](#) | [.html](#)]

We propose a novel approach to the study of Internet topology in which we use an optimization framework to model the mechanisms driving incremental growth. While previous methods of topology generation have focused on explicit replication of statistical properties, such as node hierarchies and node degree distributions, our approach addresses the economic tradeoffs, such as cost and performance, and the technical constraints faced by a single ISP in its network design. By investigating plausible objectives and constraints in the design of actual networks, observed network properties such as certain hierarchical structures and node degree distributions can be expected to be the natural by-product of an approximately optimal solution chosen by network designers and operators. In short, we advocate here essentially an approach to network topology design, modeling, and generation that is based on the concept of Highly Optimized Tolerance (HOT). In contrast with purely descriptive topology modeling, this opens up new areas of research that focus on the causal forces at work in network design and aim at identifying the economic and technical drivers responsible for the observed large-scale network behavior. As a result, the proposed approach should have significantly more predictive power than currently pursued efforts and should provide a scientific foundation for the investigation of other important problems, such as pricing, peering, or the dynamics of routing protocols.

- [4] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509-512, 1999. [[bib](#) | [.html](#)]

Systems as diverse as genetic networks or the World Wide Web are best described as networks with complex topology. A common property of many large networks is that the vertex connectivities follow a scale-free power-law distribution. This feature was found to be a consequence of two generic mechanisms: (i) networks expand continuously by the addition of new vertices, and (ii) new vertices attach preferentially to sites that are already well connected. A model based on these two ingredients reproduces the observed stationary scale-free distributions, which indicates that the development of large networks is governed by robust self-organizing phenomena that go beyond the particulars of the individual systems.

- [5] A.-L. Barabási, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173-187, 1999. [[bib](#) | [.html](#)]

Random networks with complex topology are common in Nature, describing systems as diverse as the world wide web or social and business networks. Recently, it has been demonstrated that most large networks for which topological information is available display scale-free features. Here we study the scaling properties of the recently introduced scale-free model, that can account for the observed power-law distribution of the connectivities. We develop a mean-field method to predict the growth dynamics of the individual vertices, and use this to calculate analytically the connectivity distribution and the scaling exponents. The mean-field method can be used to address the properties of two variants of the scale-free model, that do not display power-law scaling.

- [6] J. M. Carlson and J. Doyle. Highly optimized tolerance: a mechanism for power laws in designed systems. *Physical Review E*, 60(2):1412-1427, 1999. [[bib](#) | [http](#)]

We introduce a mechanism for generating power law distributions, referred to as *highly optimized tolerance* (HOT), which is motivated by biological organisms and advanced engineering technologies. Our focus is on systems which are optimized, either through natural selection or engineering design, to provide robust performance despite uncertain environments. We suggest that power laws in these systems are due to tradeoffs between yield, cost of resources, and tolerance to risks. These tradeoffs lead to highly optimized designs that allow for occasional large events. We investigate the mechanism in the context of percolation and sand pile models in order to emphasize the sharp contrasts between HOT and self organized criticality (SOC), which has been widely suggested as the origin for power laws in complex systems. Like SOC, HOT produces power laws. However, compared to SOC, HOT states exist for densities which are higher than the critical density, and the power laws are not restricted to special values of the density. The characteristic features of HOT systems include: (1) high efficiency, performance, and robustness to designed-for uncertainties, (2) hypersensitivity to design flaws and unanticipated perturbations, (3) nongeneric, specialized, structured configurations, and (4) power laws. The first three of these are in contrast to the traditional hallmarks of criticality, and are obtained by simply adding the element of design to percolation and sand pile models, which completely changes their characteristics.

- [7] J. M. Carlson and J. Doyle. Highly optimized tolerance: Robustness and design in complex systems. *Physical Review Letters*, 84(11):2529-2532, Mar. 2000. [[bib](#) | [http](#)]

We introduce HOT, a mechanism that connects evolving structure and power laws in interconnected systems. HOT systems arise, e.g. in biology and engineering, where design and evolution create complex systems sharing common features, including (1) high efficiency, performance, and robustness to designed-for uncertainties, (2) hypersensitivity to design flaws and unanticipated perturbations, (3) nongeneric, specialized, structured configurations, and (4) power laws. We introduce HOT states in the context of percolation, and contrast properties of the high density HOT states with random configurations near the critical point. While both cases exhibit power laws, only HOT states display properties (1-3) associated with design and evolution.

- [8] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited. In *Proceedings of IEEE INFOCOM*, June 2002. [[bib](#) | [.html](#)]

In a recent paper, Faloutsos et al. [[13](#)] found that the inter Autonomous System (AS) topology exhibits a power-law vertex degree distribution. This result was quite unexpected in the networking community and stirred significant interest in exploring the possible causes of this phenomenon. The work of Barabasi and Albert [[4](#)] and its application to network topology generation in the work of Medina et al. [[17](#)] have explored a promising class of models that yield strict power-law vertex degree distributions. In this paper, we reexamine the BGP measurements that form the basis for the results reported in [[13](#)]. We find that by their very nature (i.e., being strictly BGP-based), the data provides a very incomplete picture of Internet connectivity at the AS level. The AS connectivity maps constructed from this data (the original maps) typically miss 20-50% or even more of the physical links in AS maps constructed using additional sources (the extended maps). Subsequently, we find that while the vertex degree distributions resulting from the extended maps are heavy-tailed, they deviate significantly from a strict power law. Finally, we show that available historical data does not support the connectivity-based dynamics assumed in [[4](#)]. Together, our results suggest that the Internet topology at the AS level may well have developed over time following a very different set of growth processes than those proposed in [[4](#)].

- [9] Z. Dezsö and A.-L. Barabási. Halting viruses in scale-free networks. *Physical Review E*, 65, 2002. 055103. [[bib](#) | [.html](#)]

The vanishing epidemic threshold for viruses spreading on scale-free networks indicates that traditional methods, aiming to decrease a virus' spreading rate cannot succeed in eradicating an epidemic. We demonstrate that policies that discriminate between the nodes, curing mostly the highly connected nodes, can restore a finite epidemic threshold and potentially eradicate a virus. We find that the more biased a policy is towards the hubs, the more chance it has to bring the epidemic threshold above the

virus' spreading rate. Furthermore, such biased policies are more cost effective, requiring less cures to eradicate the virus.

- [10] V. M. Eguíluz, E. Hernández-García, O. Piro, and K. Klemm. Effective dimensions and percolation in hierarchically structured scale-free networks. submitted to *Physical Review E*, 2003. [[bib](#) | [http](#)]

We introduce appropriate definitions of dimensions in order to characterize the fractal properties of complex networks. We compute these dimensions in a hierarchically structured network of particular interest. In spite of the nontrivial character of this network that displays scale-free connectivity among other features, it turns out to be approximately one-dimensional. The dimensional characterization is in agreement with the results on statistics of site percolation and other dynamical processes implemented on such a network.

- [11] V. M. Eguíluz and K. Klemm. Epidemic threshold in structured scale-free networks. *Physical Review Letters*, 89(10), Sept. 2002. 108701. [[bib](#) | [http](#)]

We analyze the spreading of viruses in scale-free networks with high clustering and degree correlations, as found in the Internet graph. For the Suscetible-Infected-Susceptible model of epidemics the prevalence undergoes a phase transition at a finite threshold of the transmission probability. Comparing with the absence of a finite threshold in networks with purely random wiring, our result suggests that high clustering and degree correlations protect scale-free networks against the spreading of viruses. We introduce and verify a quantitative description of the epidemic threshold based on the connectivity of the neighborhoods of the hubs.

- [12] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the internet (extended abstract). In *Proceedings of ICALP*, pages 110-122. Springer, 2002. [[bib](#) | [.html](#)]

We give a plausible explanation of the power law distributions of degrees observed in the graphs arising in the Internet topology [13] based on a toy model of Internet growth in which two objectives are optimized simultaneously: "last mile" connection cost, and transmission delays measured in hops. We also point out a similar phenomenon, anticipated in [6], in the distribution of file sizes. Our results seem to suggest that power laws tend to arise as a result of complex, multi-objective optimization.

- [13] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proceedings of ACM SIGCOMM*, Sept. 1999. [[bib](#) | [.html](#)]

Despite the apparent randomness of the Internet, we discover some surprisingly simple power-laws of the Internet topology. These power-laws hold for three snapshots of the Internet, between November 1997 and December 1998, despite a 45% growth of its size during that period. We show that our power-laws fit the real data very well resulting in correlation coefficients of 96% or higher.

Our observations provide a novel perspective of the structure of the Internet. The power-laws describe concisely skewed distributions of graph properties such as the node outdegree. In addition, these power-laws can be used to estimate important parameters such as the average neighborhood size, and facilitate the design and the performance analysis of protocols. Furthermore, we can use them to generate and select realistic topologies for simulation purposes.

- [14] K. Klemm and V. M. Eguíluz. Growing scale-free networks with small world behavior. *Physical Review E*, 65, 2002. 057102. [[bib](#) | [http](#)]

In the context of growing networks, we introduce a simple dynamical model that unifies the generic features of real networks: scale-free distribution of degree and the small world effect. While the average shortest path length increases logarithmically as in random networks, the clustering coefficient assumes a large value independent of system size. We derive expressions for the clustering coefficient in two limiting cases: random ($C \sim (\ln N)^2 / N$) and highly clustered ($C = 5/6$) scale-free networks.

- [15] K. Klemm and V. M. Eguíluz. Highly clustered scale-free networks. *Physical Review E*, 65, Dec. 2002. 036123. [[bib](#) | [http](#)]

We propose a model for growing networks based on a finite memory of the nodes. The model shows stylized features of real-world networks: power law distribution of degree, linear preferential attachment of new links and a negative correlation between the age of a node and its link attachment rate. Notably, the degree distribution is conserved even though only the most recently grown part of the network is considered. This feature is relevant because real-world networks truncated in the same way exhibit a power-law distribution in the degree. As the network grows, the clustering reaches an asymptotic value larger than for regular lattices of the same average connectivity. These high-clustering scale-free networks indicate that memory effects could be crucial for a correct description of the dynamics of growing networks.

- [16] J. Leveille. Epidemic spreading in technological networks. Technical Report HPL-2002-287, HP Laboratories Bristol, Oct. 2002. [[bib](#) | [.html](#)]

Recent computer worms pose a major threat to large computer networks, and it is a general belief that understanding their means of propagation will help to devise efficient control strategies. This dissertation proposes a new epidemiological model to account for particular characteristics of computer worm epidemics. This new model, termed the Progressive Susceptible- Infected-Detected-Removed (PSIDR) epidemiological model, incorporates new aspects related to the availability of antivirus signatures, to the existence of direct immunization, and to the presence of a curing phase. Various costs are incorporated in the model, which allow us to determine the best strategies to fight worms. The model undergoes an extensive series of validation tests, its properties being evaluated mostly numerically. The model shows good agreement with empirical data. The paper then investigates current response strategies as well as the effect of virus throttling. The model yields both practical recommendations and new insights about the observed low prevalence of worms over the Internet.

- [17] A. Medina and I. Matta. Brite: A flexible generator of internet topologies. Technical Report BU-CS-TR-2000-005, Boston University, Boston, MA, 2000. [[bib](#)]

- [18] M. E. J. Newman. Random graphs as models of networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook of Graphs and Networks*. Wiley-VCH, Berlin, 2002. To appear. [[bib](#) | [http](#)]

The random graph of Erdos and Renyi is one of the oldest and best studied models of a network, and possesses the considerable advantage of being exactly solvable for many of its average properties. However, as a model of real-world networks such as the Internet, social networks or biological networks it leaves a lot to be desired. In particular, it differs from real networks in two crucial ways: it lacks network clustering or transitivity, and it has an unrealistic Poissonian degree distribution. In this paper we review some recent work on generalizations of the random graph aimed at correcting these shortcomings. We describe generalized random graph models of both directed and undirected networks that incorporate arbitrary non-Poisson degree distributions, and extensions of these models that incorporate clustering too. We also describe two recent applications of random graph models to the problems of network robustness and of epidemics spreading on contact networks.

- [19] M. E. J. Newman, S. Forrest, and J. Balthrop. Email networks and the spread of computer viruses. *Physical Review E*, 66, 2002. 035101. [[bib](#) | [http](#)]

Many computer viruses spread via electronic mail, making use of computer users email address books as a source for email addresses of new victims. These address books form a directed social network of connections between individuals over which the virus spreads. Here we investigate empirically the structure of this network using data drawn from a large computer installation, and discuss the implications of this structure for the understanding and prevention of computer virus epidemics.

- [20] M. E. J. Newman, M. Girvan, and J. D. Farmer. Optimal design, robustness, and risk aversion. *Phys. Rev. Lett.*, 89, July 2002. 028301. [[bib](#) | [http](#)]

Highly optimized tolerance is a model of optimization in engineered systems, which gives rise to

power-law distributions of failure events in such systems. The archetypal example is the highly optimized forest fire model. Here we give an analytic solution for this model which explains the origin of the power laws. We also generalize the model to incorporate risk aversion, which results in truncation of the tails of the power law so that the probability of disastrously large events is dramatically lowered, giving the system more robustness.

- [21] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86:3200-3203, 2001. [[bib](#) | [http](#)]

The Internet, as well as many other networks, has a very complex connectivity recently modeled by the class of scale-free networks. This feature, which appears to be very efficient for a communications network, favors at the same time the spreading of computer viruses. We analyze real data from computer virus infections and find the average lifetime and prevalence of viral strains on the Internet. We define a dynamical model for the spreading of infections on scale-free networks, finding the absence of an epidemic threshold and its associated critical behavior. This new epidemiological framework rationalize data of computer viruses and could help in the understanding of other spreading phenomena on communication and social networks.

- [22] R. Pastor-Satorras and A. Vespignani. Epidemic dynamics in finite size scale-free networks. *Physcal Review E*, 65, 2002. 035108. [[bib](#) | [http](#)]

Many real networks present a bounded scale-free behavior with a connectivity cut-off due to physical constraints or a finite network size. We study epidemic dynamics in bounded scale-free networks with soft and hard connectivity cut-offs. The finite size effects introduced by the cut-off induce an epidemic threshold that approaches zero at increasing sizes. The induced epidemic threshold is very small even at a relatively small cut-off, showing that the neglect of connectivity fluctuations in bounded scale-free networks leads to a strong over-estimation of the epidemic threshold. We provide the expression for the infection prevalence and discuss its finite size corrections. The present work shows that the highly heterogeneous nature of scale-free networks does not allow the use of homogeneous approximations even for systems of a relatively small number of nodes.

- [23] R. Pastor-Satorras and A. Vespignani. *Handbook of Graphs and Networks: From the Genome to the Internet*, chapter Epidemics and immunization in scale-free networks. Wiley-VCH, Berlin, May 2002. [[bib](#) | [http](#)]

In this chapter we want to provide a review of the main results obtained in the modeling of epidemic spreading in scale-free networks. In particular, we want to show the different epidemiological framework originated by the lack of any epidemic threshold and how this feature is rooted in the extreme heterogeneity of the scale-free networks' connectivity pattern.

- [24] C. Robert, J. M. Carlson, and J. Doyle. Highly optimized tolerance in epidemic models incorporating local optimization and regrowth. *Physical Review E*, 63, 2001. 056122. [[bib](#) | [http](#)]

In the context of a coupled map model of population dynamics, which includes the rapid spread of fatal epidemics, we investigate the consequences of two new features in Highly Optimized Tolerance (HOT), a mechanism which describes how complexity arises in systems which are optimized for robust performance in the presence of a harsh external environment. Specifically, we (1) contrast global and local optimization criteria and (2) investigate the effects of time dependent regrowth. We find that both local and global optimization lead to HOT states, which may differ in their specific layouts, but share many qualitative features. Time dependent regrowth leads to HOT states which deviate from the optimal configurations in the corresponding static models in order to protect the system from slow (or impossible) regrowth which follows the largest losses and extinctions. While the associated map can exhibit complex, chaotic solutions, HOT states are confined to relatively simple dynamical regimes.

- [25] M. Stubna and J. Fowler. An application of the highly optimized tolerance model to electrical blackouts. *International Journal of Bifurcation and Chaos*, 13(1):237-242, 2003. [[bib](#) | [.html](#) | [.pdf](#)]

The recently proposed *Highly Optimized Tolerance* (H.O.T.) model [6], [7], which aims to describe the statistics of robust complex systems in uncertain environments, is compared with data from the Western United States (W.S.C.C.) power distribution system. We use for comparison a 15-year record of all power outages occurring on the grid, measured in the size of megawatts lost and the number of customers without service. In applying the model to the power grid data, we find that the problem of determining how the resources in the system scale with event size is nontrivial given the assumptions of the model and the information about how the power grid actually operates. Further, we observe that the model agrees closely with the W.S.C.C. data for the megawatts but not the customers, and consequently propose that the assumption in the model of optimal resource distribution is not valid in general when more than one measure of event size is used. A modified H.O.T. model which allows for *resource misallocation* is introduced and we find that this model can be made to fit both data sets reasonably well.

[26] X. Zhu, J. Yu, and J. Doyle. Heavy tails, generalized coding, and optimal web layout. In *Proceedings of IEEE INFOCOM*, pages 1617-1626, Apr. 2001. [[bib](#) | [.html](#)]

This paper considers Web layout design in the spirit of source coding for data compression and rate distortion theory, with the aim of minimizing the average size of files downloaded during Web browsing sessions. The novel aspect here is that the object of design is layout rather than codeword selection, and is subject to navigability constraints. This produces statistics for file transfers that are heavy tailed, completely unlike standard Shannon theory, and provides a natural and plausible explanation for the origin of observed power laws in Web traffic. We introduce a series of theoretical and simulation models for optimal Web layout design with varying levels of analytic tractability and realism with respect to modeling of structure, hyperlinks, and user behavior. All models produce power laws which are striking both for their consistency with each other and with observed data, and their robustness to modeling assumptions. These results suggest that heavy tails are a permanent and ubiquitous feature of Internet traffic, and not an artifice of current applications or user behavior. They also suggest new ways of thinking about protocol design that combines insights from information and control theory with traditional networking.