Albert-László Barabási

Science of Networks

From Society to the Web

Since the National Science Foundation gave up its stewardship over the internet in 1995, the network appears to be living a life of its own. Routers and lines are added continuously by thousands of small and large companies, none of which are obliged to report about their activity. This uncontrolled and decentralized growth turned designers into explorers. Indeed, while until recently all internet-related research has concentrated on designing better communication protocols, lately an increasing number of scientists have begun to ask an unexpected question: What exactly did we create? One thing is clear: while entirely of human design, the web appears to have more in common with a cell or an ecological system than with a meticulously designed Swiss watch. Many diverse components, each performing a specialized job, contribute to a system that evolves and changes with an incredible speed. And we are increasingly realizing that the lack of understanding of the internet and the WWW evolution is not a computer science question, but is rooted in the absence of a scientific framework to characterize the topology of the network behind it.

Networks are everywhere. The brain is a network of nerve cells connected by axons, and cells themselves are networks of molecules connected by biochemical reactions. Societies, too, are networks of people linked by friendship, family relationships, and professional ties. On a larger scale, food webs and ecosystems can be represented as networks of species. And networks pervade technology: the internet, power grids, and transportation systems are but a few examples. Even the language we are using to convey our thoughts is itself a network of words connected by syntactic relationships.

Yet despite the importance and pervasiveness of networks, scientists have had little understanding of their structure and properties. How do the interactions of several malfunctioning genes in a complex genetic network result in cancer? How does diffusion occur so rapidly through certain social and communications networks, leading to epidemics of diseases and computer viruses like the Love Bug? How do some networks continue to function even after the vast majority of their nodes have failed?

Recent research has begun to answer such questions.[1] Over the past few years, scientists from a variety of fields have discovered that complex networks seem to have an underlying architecture that is guided by universal principles. We have found, for instance, that many networks – from the World Wide Web to the cell's metabolic system to the actors in Hollywood – are dominated by a relatively small number of nodes that are highly connected to other nodes. These important nodes, called "hubs", can greatly affect a network's overall behaviour, for instance, making it remarkably robust against accidental failures but extremely vulnerable to coordinated attacks.

The Random Network Paradigm

For over 40 years science treated all complex networks as being completely random. This paradigm has its roots in the work of two Hungarian mathematicians, Paul Erdős and Alfréd Rényi, who in 1959, aiming to describe networks seen in communications and life sciences, suggested that we should build networks randomly.[2] Their recipe was simple: take N nodes and connect them by L randomly placed links. The simplicity of the model and the elegance of some of the related theorems proposed by Erdős and Rényi have revitalized graph theory, leading to the emergence of a new field in mathematics, focusing on random networks.[3]

An important prediction of random network theory is that despite the fact that we place the links randomly, the resulting network is deeply democratic, as most nodes have approximately the same number of links. Indeed, in a random network the nodes follow a Poisson distribution with a bell shape, and it is extremely rare to find nodes that have significantly more or fewer links than the average node. Random networks are also called exponential networks because the probability that a node is connected to k other nodes decreases exponentially for large k (Fig. 1). But the Erdős-Rényi model raises an important question: do we believe

[1] A.-L. Barabási, *Linked: The New Science of Networks*, Cambridge, MA: Perseus Publishing, 2002; R. Albert and A.-L. Barabási, "Statistical Mechanics of Complex Networks", *Review of Modern Physics* 74 (2002), pp. 47–97; R. Pastor-Satorras and A. Vespignani, *Evolution and Structure of the Internet: A Statistical Physics Approach*, Cambridge University Press, 2004; A.-L. Barabási and Z. N. Oltvai, "Network Biology: Understanding the Cells's Functional Organization", *Nature Reviews Genetics* 5 (2004), pp. 101–113.

[2] P. Erdős and A. Rényi, "On Random Graphs I", *Publ. Math.* (Debrecen) 6 (1959), pp. 290–297; P. Erdős and A. Rényi, "On the Evolution of Random Graphs", *Publ. Math. Inst. Hung. Acad. Sci.* 5 (1960), pp. 17–60.

[3] B. Bollobás, *Random Graphs*, London: Academic Press, 1985.

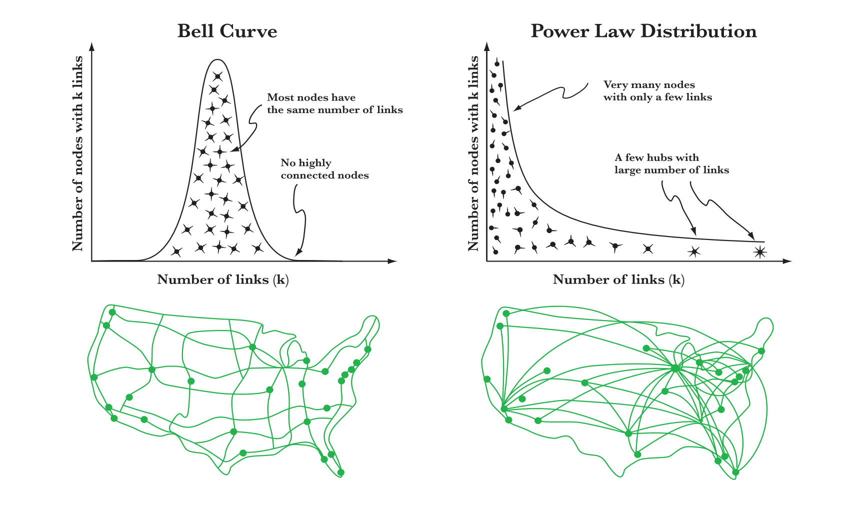

Figure 1
Random and scale-free networks

The degree distribution of a random network follows a Poisson curve very close in shape to the Bell Curve, telling us that most nodes have the same number of links, and nodes with a very large number of links don't exist (top left). Thus a random network is similar to a national highway network, in which the nodes are the cities, and the links are the major highways connecting them. Indeed, most cities are served by roughly the same number of highways (bottom left). In contrast, the power law degree distribution of a scale-free network predicts that most nodes have only a few links, held together by a few highly connected hubs (top right). Visually this is very similar to the air traffic system, in which a large number of small airports are connected to each other via a few major hubs (bottom right).

(After A.-L. Barabási, *Linked*)

that networks observed in nature are truly random? Could the internet offer us the relatively fast and seamless service if the computers were randomly connected to each other? Or, to take a more extreme example, would you be able to read this article if in a certain moment the chemicals in your body would decide to react randomly to each other, bypassing the rigid chemical web they normally obey? Intuitively the answer is no – we all feel that behind each complex system there is an underlying network with non-random topology. Thus our challenge is to unearth the signatures of order from this apparent chaos of millions of nodes and links.

The WWW contains over a billion documents, which represent the nodes of this complex web. They are connected by URLs that allow us to navigate from one document to another (Fig. 2a). To analyze its prop-

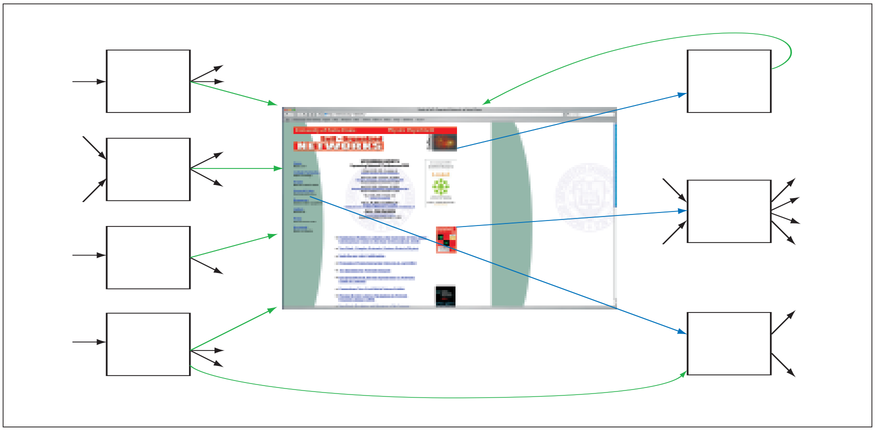

Figure 2a

The nodes of the World Wide Web are web documents, each of which is identified by an unique uniform resource locator, or URL. Most documents contain URLs that link to other pages. These URLs represent outgoing links, three of which are shown (blue arrows). Currently there are about 80 documents worldwide that point to our website www.nd.edu/~networks, represented by the incoming green arrows. While we have complete control over the number of the outgoing links, k_{out}*, from our webpage, the number of incoming links,* k_{in}*, is decided by other people, and thus characterizes the popularity of the page.*

(After A.-L. Barabási, "The Physics of the Web", *Physics World*, 2001, pp. 33–38)

erties, we need to obtain a map, telling us how the pages link to each other. This information is routinely collected by search engines, such as Google or AltaVista, but they are often reluctant to share it for research purposes. Thus we needed to obtain a map of our own. This is exactly what we did in 1998 – we wrote a robot or web crawler that, starting from a given webpage, collected all the outgoing links, and followed those links to visit more pages and collect even more links.[4] Through this iter-

[4] R. Albert, H. Jeong and A.-L. Barabási, "Diameter of the World Wide Web", *Nature* 401 (1999), pp. 130–131.

ative process we mapped out a small fraction of the web.

As the WWW is a directed network, each document can be characterized by the number of outgoing (k_{out}) and incoming (k_{in}) links. The first quantity that we investigated was the outgoing (incoming) degree distribution, which represents the probability *P(k)* that a randomly selected webpage has exactly k_{out} (k_{in}) links. Guided by random graph theory, we expected that *P(k)* would follow a Poisson distribution. Thus it was rather surprising when the data indicated that *P(k)* decayed following a power law (Fig. 2b),

$$P(k) \sim k^{\gamma}, \tag{1}$$

where $\gamma_{out} \cong 2.45$ ($\gamma_{in} \cong 2.1$).

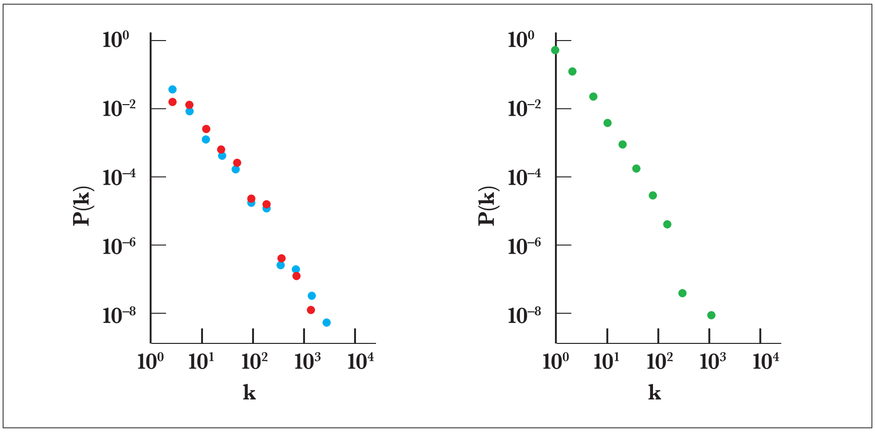

Figure 2b

Left: *The probability that a webpage has* k_{in} *(blue) or* k_{out} *(red) links follows a power law. The results are based on a sample of over 325 000 webpages collected by Hawoong Jeong*. Right: *The degree distribution of the internet at the router level, where k denotes the number of links a router has to other routers. This research by Ramesh Govindan from University of Southern California is based on over 260,000 routers and demonstrates that the internet exhibits power-law behaviour.*

(After A.-L. Barabási, "The Physics of the Web")

There are major topological differences between networks with Poisson or power-law connectivity distribution (as illustrated in Fig. 1). Indeed, for random networks most nodes have approximately the same number of links, $k \approx \langle k \rangle$, where $\langle k \rangle$ represents the average degree. The exponential decay of $P(k)$ guarantees the absence of nodes with significantly more links than $\langle k \rangle$. In contrast, the power-law distribution implies that nodes with only a few links are abundant, but a small, negligible minority have a very large number of links. For example, the highway map, with cities as nodes and highways as links, is an exponential network, as most cities are at the intersection of two to five highways. On the other hand, a scale-free network looks more like an airline routing map, displayed routinely in glossy flight magazines: most airports are served only by a few carriers, but there are a few hubs, such as Chicago or Frankfurt, from which links emerge to almost all other U.S. or European airports, respectively. Thus, just as the smaller airports, in the WWW the majority of the documents have only a few links. These few links are not sufficient to ensure that the network is fully connected, a function guaranteed by a few highly connected hubs, that hold the nodes together.

While for a Poisson distribution a typical node has $k \cong \langle k \rangle$ links, the average, $\langle k \rangle$, is not particularly significant for a power-law distribution. This absence of an intrinsic scale in k prompted us to name networks with power-law degree distribution scale-free.[5] The finding that the WWW is a scale-free network raised an important question: would such inhomogenous topology emerge in other complex systems as well? Recently answer to this question came from an unexpected direction: the internet. The internet forms a physical network, whose nodes are routers and domains, while links represent the various physical lines, such as phone lines, optical cables, that connect them (Fig. 2c). Due to the physical nature of the connections, this network was expected to be different from the WWW, where adding a link to an arbitrary remote page is as easy as linking to a computer in the next room. To the surprise of many, the network behind the internet also appears to follow a power law degree distribution. This was first noticed by three brothers, Michalis, Petros, and Christos Faloutsos, computer scientists at U.S. and Canadian universities, who analyzed the internet at the router and domain level (see Fig. 2c). In each of these cases they found that the degree distribution follows a power law with an exponent γ=2.5 for the router network and γ=2.2 for the domain map, indicating that the wiring of the internet is also dominated by several highly connected hubs.[6]

---

[5] A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks", *Science* 286 (1999), pp. 509–512.

[6] M. Faloutsos, P. Faloutsos and C. Faloutsos, "On Power-Law Relationships of the Internet Topology", ACM SIGCOMM 99, *Comput. Commun. Rev.* 29 (1999), p. 251.

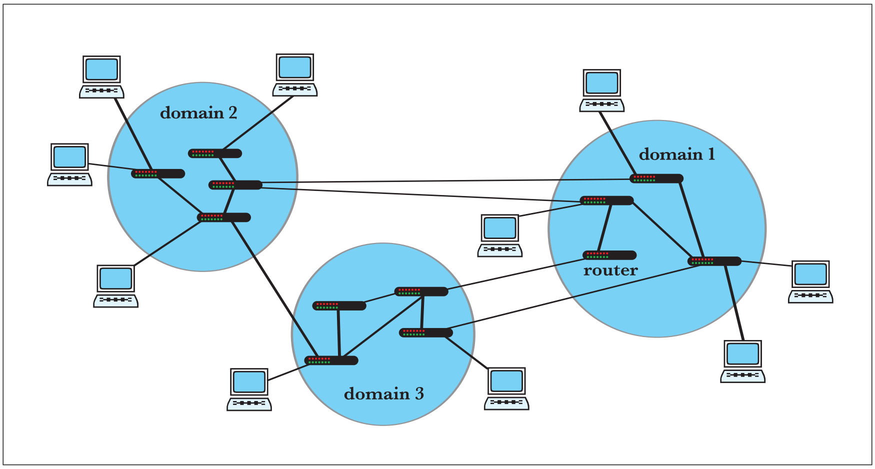

Figure 2c

The internet is a network of routers that navigate packets of data from one computer to another. The routers are connected to each other by various physical or wireless links and are grouped into several domains.

(After A.-L. Barabási, "The Physics of the Web")

Nineteen Degrees of Separation

Stanley Milgram, a Harvard sociologist, surprised the world in 1967 with a bold claim: in society two people are typically five to six handshakes away from each other.[7] That is, despite the six billion inhabitants of our planet, we live in a "small-world". This feature of social networks came to be known as "six degrees of separation" after John Guare's brilliant Broadway play and movie.[8] On top of this, sociologists have repeatedly argued that nodes in social networks are grouped in small clusters, representing circles of friends and acquaintances, in which each node is connected to all other nodes, with only sparse links to the outside world.[9] While the existence of such local clustering and small world behaviour agrees with our intuition, these features were not expected to be relevant beyond social systems. The question is: does the internet and the WWW

[7] S. Milgram, "The Small World Problem", *Psychology Today* 1 (1967), pp. 60–67.

[8] J. Guare, *Six Degrees of Separation*, New York: Vintage Books, 1990.

[9] M. S. Granovetter, "The Strength of Weak Ties", *American Journal of Sociology* 78 (1973), pp. 1360–1380.

follow this paradigm? For a proper answer we need a full map of the web. But, as Steve Lawrence and Lee Giles have shown in 1998, even the largest search engines cover only 16% of the web.[10] This is where the tools of statistical mechanics come handy – we need to infer the properties of the full web from a finite sample. To achieve this we constructed small models of the WWW in the computer, making sure that the link distribution matches the measured functional form.[11] Next we identified the shortest distance between two nodes, and averaged this over all pairs of nodes, obtaining the average node separation, d. Repeating this for networks of different sizes, using finite size scaling, a standard procedure in statistical mechanics, we inferred that d depends on the number of nodes, N, as $d = 0.35 + 2.06 \log(N)$. As in 1999 the WWW had 800 million nodes, this expression predicted that the typical shortest path between two randomly selected pages is around 19 – assuming that there is such a path, which is not guaranteed thanks to the web's directed nature. An extensive study by an IBM-Compaq-AltaVista collaboration has consequently found that for the 200 million nodes this distance is 16 – not too far from 17 predicted by our formula for a sample of this size.[12] These results clearly indicated that the WWW represents a small world, i.e. the typical number of clicks between two webpages is around 19, despite the over billion pages out there. And as Lada Adamic from Stanford University has shown, the WWW displays a high degree of clustering as well, the probability that two neighbours of a given node are linked together being much larger than the value expected for a clustering-free random network.[13] Results from our group indicated that the internet follows suit – the typical separation between two routers is 9, i.e. a package can reach any router within ten hops,[14] and the network is highly clustered, demonstrating that the small-world paradigm has rapidly infiltrated the Earth's newly developing electronic skin as well.

[10] S. Lawrence and C. L. Giles, "Searching the World Wide Web", *Science* 280 (1998), pp. 98–100; S. Lawrence and C. L. Giles, "Accessibility of Information on the Web", *Nature* 400 (1999), pp. 107–109.

[11] R. Albert, H. Jeong and A.-L. Barabási, "Diameter of the World Wide Web", cf. note 4 above.

[12] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajalopagan, R. Stata, A. Tomkins and J. Weiner, "Graph Structure in the Web", *Comput. Netw.* 33 (2000), p. 309.

[13] L. A. Adamic and B. A. Huberman, "Growth Dynamics of the World-Wide Web", *Nature* 401 (1999), p. 131.

[14] S.-H. Yook, H. Jeong, and A.-L. Barabási, "Modeling the Internet's Large-Scale Topology", *Proceedings of the National Academy of Sciences* 99 (2002), pp. 13382–13386.

Evolving Networks

Why do so different systems as the physical network of the internet or the virtual web of the WWW develop similar scale-free networks? We have recently traced back the emergence of the power law degree distribution to two common mechanisms, absent from the classical graph models, but present in many complex networks.[15] First, traditional graph theoretic models assumed that the number of nodes in a network is fixed. In contrast, the WWW continuously expands by the addition of new webpages, or the internet grows by the installation of new routers and communication links. Second, while random graph models assume that the links are distributed randomly, most real networks exhibit preferential attachment: there is a higher probability to connect to a node with a large number of links. Indeed, we link our webpage with higher probability to the more connected documents on the WWW, as these are the ones we know about; network engineers tend to connect their institution to the internet through points where there is high bandwidth, which inevitably implies a high number of other consumers, or links. Based on these two ingredients, we constructed a simple model in which a new node is added at every timestep to the network, linking to some of the nodes present in the system (Fig. 3).[16] The probability $\Pi(k)$ that a new node connects to a node with k links follows preferential attachment, i.e.

$$\Pi(k) = \frac{k}{\Sigma_i k_i} \tag{2}$$

where the sum goes over all nodes. Numerical simulations indicate that the resulting network is indeed scale-free, the probability that a node has k links following (1) with exponent $\gamma = 3$.[17] This simple model illustrates how growth and preferential attachment jointly lead to the appearance of the hub hierarchy: a node rich in links will increase its connectivity faster than the rest of the nodes, since incoming nodes link to it with higher probability, a rich-gets-richer phenomenon present in many competitive systems.

[15] A.-L. Barabási and R. Albert, “Emergence of Scaling in Random Networks”, cf. note 5 above.

[16] *Ibid.*

[17] A.-L. Barabási, R. Albert and H. Jeong, “Mean-Field Theory for Scale-Free Random Networks”, *Physica* A 272 (1999), pp. 173–187.

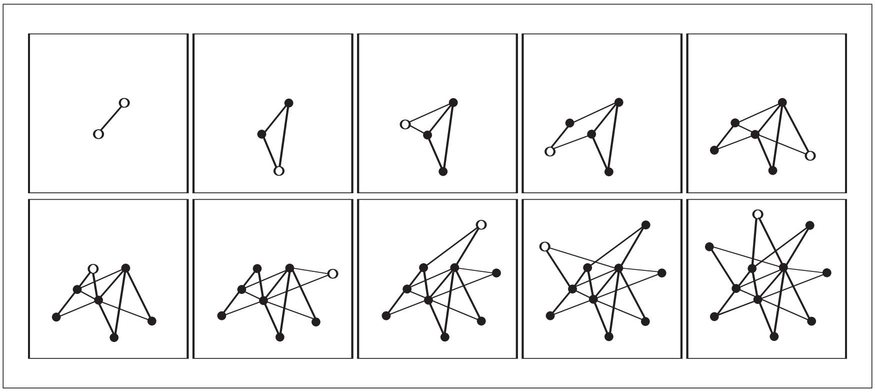

Figure 3
The birth of a scale-free network

The scale-free topology is a natural consequence of the ever-expanding nature of real networks. Starting from two connected nodes (top left), in each panel a new node (shown as an empty circle) is added to the network. When deciding where to link, new nodes prefer to attach to the more connected nodes. Thanks to growth and preferential attachment, a few highly connected hubs emerge.

(After A.-L. Barabási, *Linked*)

Networks traditionally were viewed as static objects, the role of the modeler being to find a way to place the links between a constant number of nodes such that the resulting network looks similar to the network we wish to model. In contrast, the scale-free model views networks as dynamical systems, incorporating the fact that they self-assemble and evolve in time through the addition and removal of nodes and links. Such a dynamical approach follows the long tradition of physics-based modeling, aiming to capture what nature did when it assembled these networks. The expectation behind these modeling efforts is that if we capture the microscopic processes that drive the placement of links and nodes, the structural elements and the topology will follow from these. In addition, viewing evolving networks as dynamical systems allows us to predict many of their properties analytically.

Bose-Einstein Condensation

In most complex systems nodes vary in their ability to compete for links. For example, some webpages, through a mix of good content and marketing, acquire a large number of links in a short time, easily passing less popular sites that have been around much longer. A good example is the Google search engine: a relatively latecomer with an excellent product in less than two years became one of the most connected nodes of the WWW. This competition for links can be incorporated into the scale-free model by adding to each node a fitness, describing its ability to compete for links at the expense of other nodes.[18] Assigning a randomly chosen fitness η_i to each node i modifies the growth rate in (3) to

$$\Pi(k_i) = \frac{\eta_i k_i}{\Sigma_j \eta_j k_j} \tag{3}$$

The competition generated by the unequal fitness leads to multiscaling: the connectivity of a given node follow $k_i(t) \cong t^{\beta(\eta)}$, where $\beta(\eta)$ increases with η, allowing fitter nodes with large η to join the network at some later time and overcome the older but less fit nodes.

The competitive fitness models appear to have close ties to Bose-Einstein condensation, currently one of the most investigated problems in condensed matter physics. Indeed, we have recently found[19] that the fitness model can be mapped exactly into a Bose gas by replacing each node with an energy level of energy $\varepsilon_i = e^{-\beta\eta_i}$. According to this mapping, links connected to node i are replaced by particles on level ε_I, and the behaviour of the Bose gas is uniquely determined by the distribution $g(\varepsilon)$ from which the fitnesses are selected. One expects that the functional form of $g(\varepsilon)$ is system dependent: the attractiveness of a router for a network engineer comes from a rather different distribution that the fitness of a .com company competing for customers. For a wide class of fitness distributions a fits-gets-richer phenomena emerges, in which, while the fittest node acquires more links than its less fit counterparts, there is no clear winner. On the other hand, certain fitness distributions can result in Bose-Einstein condensation, which in the network language corresponds to a winner-takes-all phenomenon: the fittest node emerges as a clear winner, developing a condensate by acquiring a finite fraction of the links, independent of the size of the system.

---

[18] G. Bianconi and A.-L. Barabási, "Competition and Multiscaling in Evolving Networks", *Europhysics Letters* 54 (2001), pp. 436–442.

[19] G. Bianconi and A.-L. Barabási, "Bose-Einstein Condensation in Complex Networks", *Physical Review Letters* 86 (2001), pp. 5632–5635.

The Achilles' Heel of the Internet

As the world economy becomes increasingly dependent on the internet, a much voiced concern arises: can we maintain its functionality under inevitable failures or frequent hacker attacks? The good news is that so far the internet has proven rather resilient against failures: while about 3% of the routers are down at any moment, we rarely observe major disruptions. The question is, where does this robustness come from? While there is significant error tolerance built into the protocols that govern package switching, lately we are learning that the scale-free topology also plays a crucial role. In trying to understand the network component of error tolerance, we get help from percolation, a much studied field of physics. Percolation theory tells us that the random removal of nodes from a network will result in an inverse percolation transition: as a critical fraction, f_c, of the nodes is removed, the network should fragment into tiny, non-communicating islands of nodes. To our considerable surprise simulations on scale-free networks did not support this prediction:[20] we could remove as many as 80% of the nodes, and the remaining nodes still formed a compact cluster. The mystery was resolved by Shlomo Havlin of Bar-Ilan University and his collaborators, who have shown that as long as the connectivity exponent γ is smaller than 3 (which is the case for most real networks, including the internet) the critical threshold for fragmentation is $f_c=1$.[21] This is a wonderful demonstration that scale-free networks cannot be broken into pieces by the random removal of nodes. This extreme robustness to failures is rooted in the inhomogeneous network topology: as there are far more small nodes than hubs, random removal will most likely hit these. But the removal of a small node does not create a significant disruption in the network topology, just like the closure of a small local airport has little impact on international air traffic, explaining the network's robustness against random failures. The bad news is that the inhomogeneous topology has its drawbacks as well: scale-free networks are rather vulnerable to attacks.[22] Indeed, the absence of a tiny fraction of the

[20] R. Albert, H. Jeong and A.-L. Barabási, "The Internet's Achilles' Heel: Error and Attack Tolerance in Complex Networks", *Nature* 406 (2000), pp. 378–382.

[21] R. Cohen, K. Erez, D. ben-Avraham and S. Havlin, "Resilience of the Internet to Random Breakdowns", *Physical Review Letters* 85 (2000), pp. 4626–4628; R. Cohen, K. Erez, D. ben-Avraham and S. Havlin, "Breakdown of the Internet under Intentional Attack", *Physical Review Letters* 86 (2001), pp. 3682–3685.

[22] R. Albert, H. Jeong and A.-L. Barabási, "The Internet's Achilles' Heel", cf. note 20 above.

most connected nodes will break the network into pieces. These findings uncovered the underlying topological vulnerability of scale-free networks: while the internet is not expected to break under the random failure of the routers and lines, well-informed hackers can easily design a scenario to handicap the network.

Scale-Free Epidemics

Knowledge about scale-free networks also has implications for understanding the spread of computer viruses, diseases, and fads. Diffusion theories, intensively studied for decades by both epidemiologists and marketing experts, predict a critical threshold for the propagation of something throughout a population. Any virus that is less virulent (or a fad that is less contagious) than that well-defined threshold will inevitably die out, while those above the threshold will multiply exponentially, eventually penetrating the entire system.

Recently, though, Romualdo Pastor-Satorras from Unversitat Politecnica de Catalunya in Barcelona and Alessandro Vespignani from Indiana University in Bloomington reached a startling conclusion.[23] They found that on a scale-free network the threshold is zero. That is, all viruses, even those that are only weakly contagious, will spread and persist in the system. This explains why Love Bug, the most damaging virus thus far, having shut down the British Parliament in 2000, was still the seventh most frequent virus even a year after its introduction and supposed eradication. Hubs are the key to that surprising behaviour. Because hubs are highly connected, at least one of them will tend to be infected by any corrupted node. And once a hub has been infected, it will broadcast the virus to numerous other sites, eventually compromising other hubs that will then help spread the virus throughout the entire system.

Because biological viruses spread on social networks, which in many cases appear to be scale-free, this result suggests that scientists should take a second look at the volumes of research written on the interplay of network topology and epidemics. Specifically, in a scale-free contact network, the traditional public-health approach of random immunization could easily fail because it will likely neglect some of the hubs. Research in scale-free networks suggests an alternative approach: by targeting the hubs, or the most connected individuals, the immunizations would have

[23] R. Pastor-Satorras and A. Vespignani, "Epidemic Spreading in Scale-Free Networks", *Physical Review Letters* 86 (2001), pp. 3200–3203.

to reach only a small fraction of the population.[24] But identifying the hubs in a social network is much more difficult than in other types of systems like the internet. Nevertheless, Reuven Cohen and Shlomo Havlin of Bar-Ilan University in Israel, together with Daniel ben-Avraham of Clarkson University, New York, recently proposed a clever solution:[25] immunize a small fraction of the random acquaintances of randomly selected individuals – a procedure that selects hubs with high probability because they are acquainted to many people. That approach, though, raises important ethical dilemmas. For instance, even if the hubs could be identified, should they have priority for immunizations and cures?

In many business contexts, people want to start epidemics, not stop them. Many viral marketing campaigns, for instance, specifically try to target hubs to spread the adoption of a product as quickly as possible. Obviously, such a strategy is not new. As far back as the 1950s, a study funded by the pharmaceutical giant Pfizer discovered the important role that hubs play in how quickly a community of doctors will adopt a new drug. Indeed, marketers have intuitively known for some time that certain customers are much more effective in spreading promotional buzz about new products and fads. But recent work in scale-free networks provides the scientific framework and mathematical tools to probe that phenomenon more rigorously.

Outlook

Although scale-free networks are surprisingly pervasive, prominent exceptions exist. For example, the highway network and the power grid in the United States are not scale-free. With other networks, the data are inconclusive. The small size of food webs reliably mapped out, telling us how species feed on each other, did not allow scientists to reach a clear conclusion on the network's type. The absence of large-scale connectivity maps of the brain does not allow us to address the nature of this crucial network either. Most networks seen in materials science, such as the crystal lattice describing the interactions between atoms in solids, are not scale-free either, but all atoms have the same number of links to other atoms, leading to a quite regular network topology.

---

[24] Z. Dezső and A.-L. Barabási, "Can We Stop the AIDS Epidemic?", *Physical Review* E 65 (2002), 055103(pp. 1–4); R. Pastor-Satorras and A. Vespignani, "Immunization of Complex Networks", *Physical Review* E 65 (2002), 036104.

[25] R. Cohen, S. Havlin and D. ben-Avraham, "Efficient Immunization Strategies for Computer Networks and Populations", *Physical Review Letters* 91 (2003), 247901(pp. 1–4).

Perhaps even more important are the other parameters that define a network's structure. One such characteristic is the diameter, or path length, of a network: the largest number of hops required to get from any node to any other node by following the shortest route possible. Networks with short diameters are referred to as "small-world", and much research is currently investigating this and other related phenomena, such as node clustering and hierarchy.

Finally, understanding a network's structure is just part of the story. There might be steep costs, for instance, with the addition of each incremental link to a node, which could prevent certain networks like the U.S. highway system from becoming scale-free. In food chains, some prey are easier to catch than others, and that fact has a profound impact on the overall network. With social networks, links to family members are very different in strength than links to acquaintances. For transportation, transmission, and communications systems like the internet, the congestion of traffic along specific links is a major consideration.

In essence, we have studied complex networks first by ignoring the details of their individual links and nodes. By distancing ourselves from those particulars, we have been able to better glimpse some of the organizing principles behind these seemingly incomprehensible systems. At the very least, knowledge from this endeavour has led to the rethinking of many basic assumptions. In the past, for example, researchers would model the internet as a random network to test how a new routing protocol might affect system congestion. But, as we now know, the internet is a scale-free system with behaviour that is dramatically different than that of a random network. Consequently, researchers have been busy revamping the computer models they've been using to simulate the internet. Similarly, knowledge of the properties of scale-free networks will be invaluable in a number of other fields, particularly as we move beyond network topologies to probe the intricate and often subtle dynamics taking place within those complex systems.

The advances discussed here represent only the tip of the iceberg. Networks represent the architecture of the complexity. But to fully understand complex systems, we need to move beyond this architecture, and uncover the laws governing the underlying dynamical processes, such as internet traffic or reaction kinetics in the cell. Most important, we need to understand how these two layers of complexity, architecture and dynamics, evolve together. These are all formidable challenges for physicists, biologists, and mathematicians alike, inaugurating a new era that Stephen Hawking recently called the "century of complexity".