

## Network Laws

Many networks, physical and social, are complex and scale-invariant. This has important implications from the spread of epidemics and innovations to protection from attack.

**N**etworks are hot. The Internet has made it possible to observe and measure linkages

representing relationships of all kinds. We now recognize networks everywhere: air traffic, banking, chemical bonds, data communications, ecosystems, finite element grids, fractals, interstate highways, journal citations, material structures, nervous systems, oil pipelines, organizational networks, power grids, social structures, transportation, voice communication, water supply, Web URLs, and more.

Several fields are collaborating on the development of network theory, measurement, and mapping: mathematics (graph theory), sociology (networks of influence and communication), computing (Internet), and business (organizational networks). This convergence has produced useful results for risk assessment and reduction in complex infrastructure networks,

attacking and defending networks, protecting against network connectivity failures, operating busi-

nesses, spreading epidemics (pathogens as well as computer viruses), and spreading innovation. Here, I will survey the fundamental laws of networks that enable these results.

### Defining a Network

A network is usually defined as a set of nodes and links. The nodes represent entities such as persons, machines, molecules, documents, or businesses; the links represent relationships between pairs of entities. A link can be directed (one-way relationship) or undirected (mutual relationship). A hop is a transition from one node to another across a single link separating them. A path is a series of hops. Networks are very general: they can represent any kind of relation among entities.

Some common network topologies (interconnection patterns) have their own names: clique or island (a connected sub-network that may be isolated from other cliques), hierarchical network (tree structured), hub-and-spoke network (a special node, the hub, connected directly to every other node), and multi-hub network (several hubs connected directly to many nodes). Some network topologies are planned, such as the electric grid, the interstate highway system, or

---

## It should be no surprise that physical communication networks inherit connections from social networks.

---

the air traffic system; others are unplanned. In his seminal papers about the Internet, Paul Baran proposed that a planned, distributed network would be more resilient to failures than a hub-and-spoke network.

A host of physical systems easily fit a network model. Perhaps less obvious is that human social networks also fit the model. The individuals of an organization are linked by their relationships—who emails whom, who seeks advice from whom, or who influences whom. Software tools such as InFlow and Netform are used by organizations to map their social networks and identify individuals that serve critical networking roles. Three roles are prominent: the hub, a person with links to many others; the broker (or gatekeeper), a person who is the only connection between cliques; the bridge (or pulsetaker), a person who links several cliques and can see opportunities for exchange between them. Such structures are not visible in the hierarchical organization chart, which maps only the structure of authority to make declarations.

It should be no surprise that physical communication networks inherit connections from social networks. The structure of links in the Web reflects individual perceptions of the most useful information. The structure of connections among Internet routers reflects individual choices for the shortest, fastest connections to the sites with the most valuable data. This is where the convergence of social networking and physical networking is the most apparent.

### **Vast Networks**

There has been considerable interest in networks containing thousands or millions of nodes. These networks are so large and complex that it is hard to even draw accurate maps of them. The Internet is the first vast network to be mapped and carefully measured. Statistical sampling must be used because it is impossible to measure every node. Bill Cheswick has created beautiful Internet maps from sampled “trace-route” data (traceroute is a tool that returns the IP addresses on a path from a sender to any Internet address; see [research.lumeta.com/ches/map/](http://research.lumeta.com/ches/map/)).

Others collect Internet statistics by sampling routing tables and measuring the numbers of direct connections from a router to others.

These statistics yield surprising insights into basic networking questions of wide interest: What is the longest connection path in a vast network? How fast can an epidemic or an idea spread in a network? Do vast networks have hubs or are their connections more uniformly distributed? What is the effect of a random failure on the connectivity of the remaining nodes? Can an attacker splinter the network with a focused attack? How should a defender protect the network?

Vast, random networks were first studied by the mathematicians Paul Erdos and Alfred Renyi in a series of eight famous papers published between 1959 and 1968 [6]. In their model, links are randomly placed among a fixed set of nodes. The probability distribution of the number of links at a node has a bell-shaped curve centered on the mean number of links. When the average number of links is less than 1, there is a high probability that the network is a set of dis-

connected islands; but if the average is 1 or greater, the entire network is connected. This model stood for many years as the “gold standard” for random networks.

But when researchers started measuring the connection distributions of real networks they found that the actual distributions do not match those predicted by Erdos. Instead they found that the probability of  $k$  links at a node is proportional to  $(1/k)^p$ , where the power  $p$  is typically 2 to 3. This distribution is confirmed experimentally because the plot of the number of occurrences of nodes with exactly  $k$  links versus  $k$  on log-log paper yields a straight line of slope  $-p$ . It is called a power-law distribution. A power-law distribution documents a system in which the great majority of nodes have very few connections, but a very few nodes have a great many connections.

In contrast, the random model predicts that very high connectivity is exceedingly rare and unlikely to be observed at all. A power-law distribution has no humps that would identify that some degrees of connection are preferred over others. Power-law distributions have extremely large (or infinite) standard deviations, which means that no confidence can be placed in a prediction of the connectivity of any node sampled at random. The number of nodes of connectivity  $2k$  is a fixed fraction  $(1/2^p)$  of the number of connectivity  $k$ . For these reasons,

networks with power-law connection statistics are called scale-free networks.

Scale-free networks are common [1–4]. The relation between earthquakes and their Richter magnitudes is scale-free with  $p=2$ . The distribution of the numbers of Hollywood actors who starred in the same films is scale-free with  $p=2.3$ . The Web is scale-free with  $p=2.1$  to 2.4. Journal citations are scale-free with  $p=3$ . The electric power grid is scale-free with  $p=4$ . The term scale-free has also been applied to fractals, which are structures whose components are similar to the overall structure.

### Properties of Scale-Free Networks

What are the base conditions that make a network scale-free? Albert-Laszlo Barabasi says there are two:

- Growth: new nodes appear at random times; and
- Preferential attachment: a new node connects to an existing node with probability proportional to the number of connections already at that node [1–3].

In other words, well-connected nodes tend to attract more connections than poorly connected nodes. Any network meeting these two conditions will evolve into a scale-free state. The phrase “evolve into” is often stated “self-organize into” because there is no outside agent forcing the

organization. The amazing fact that so many different kinds of networks are scale free is explained by growth and preferential attachment. Why most of them have  $p$  between 2 and 3 is still an open question.

With this model, the oldest nodes will be the most connected, a condition called first market advantage. However, it is not always true that the first-comers are the most connected. Newcomer Google took over as a major hub in the Web because it offered more value than other search engines. When he modified his attachment rule so that preference depends both on an arbitrary preference constant as well as connectivity, Barabasi found the networks are still scale-free [1, 3].

Since these conditions are true of social networks, large social networks tend to be scale-free. Communication networks set up within these social networks follow the social communication patterns and tend also to be scale-free.

Scale-free networks have a small diameter (maximum distance between two nodes); the Web’s diameter is 19 [1, 2]. Some networks are not scale-free, but their diameters are still small. Small-world networks were first discovered by Stanley Milgram in his famous postal experiment in 1967, when he found that letters addressed only to a name would find their way to the recipient

# The Profession of IT |

within six hops. The phrase “six degrees of separation” comes from this work [8, 10].

## Vulnerabilities

A node failure means that any path through that node is unusable; it’s as though the node and its links were deleted from the

number of hubs. If one of them is deleted, a large number of connections go with it, and that may cause the network to fragment into many parts.

Thus an attacker’s strategy will be to locate and attack the hubs. Internet denial-of-service attacks disable a node by overloading it:

them. This can have enormous payoffs in the reliability and resilience of the network [7].

## Innovations

As I have discussed in previous columns, innovations are transformations of practice within communities [5]. How do inno-

---

## The vast majority of innovations affect only small groups and a few affect enormous numbers of people.

---

network. In some networks, a node failure might divide the network into two or more disconnected pieces—for example, on losing a broker or gatekeeper in a social network. In others, a random node failure would not affect overall connectivity—for example, the Internet “routes around damage.” The properties of scale-free networks tell us two things about damage caused by node failures:

*Good news:* A random node failure in a scale-free network has negligible effect on connectivity. This is because the vast majority of nodes have only a few connections; deleting one of them will hardly be noticed.

*Bad news:* Failure of a hub in a scale-free network can significantly damage the network’s connectivity. The power law distribution guarantees that there will be a small

number of hubs. If one of them is deleted, a large number of connections go with it, and that may cause the network to fragment into many parts.

Conversely, a defender’s strategy is to harden the hubs. In fact, a defender will not have sufficient resources to protect every node in a vast network. Therefore, protecting the hubs is the best policy.

Viruses spread fast in scale-free networks because the network diameter is small. Launch a virus at any node and within a few hops it reaches a hub; within a few more hops it reaches everyone. Some virus experts believe that if the network hubs could be secured successfully against viruses, most virus problems would disappear.

Guardians of critical infrastructure have found that the networks they protect are scale-free. The only sound policy is to identify the hubs of their networks and devote the limited resources to protecting

innovations propagate?

Innovations propagate through the social network of the community. The idea (or proposal) for a new practice passes from one member to another through their collaboration links. When it reaches a new member, the idea will either be rejected, or it will be adopted and offered to the member’s immediate colleagues. Organization theorists have studied this with the help of social network maps. Innovations are most likely to spread if the hubs adopt and recommend them; brokers and bridges play important roles in jumping the idea to new groups.

Given these dynamics, it is likely that the “size of an innovation”—the number of people ultimately involved—also follows a power law. That is, the vast majority of innovations affect only small

groups and a few affect enormous numbers of people. Because the standard deviation of a power law distribution is very large or infinite, there can be little confidence in an advance prediction of the size of an innovation.

That said, the spread of innovations is not as random as the discussion might make it seem. It has been well documented that some people are better at conceiving and spreading innovations than others. In a previous column, I summarized the personal foundational practices of skilled innovators [5]. From this we can infer a few guidelines:

*Locate the hubs of the network you're trying to influence.* Bring them on board. That will make it much easier for many others in the community to follow suit. Conversely, if you pick random individuals ("cold calls"), the chances are very high that you will pick a poorly connected node. Even if that individual adopts your practice, it will take a long time to propagate to others.

*Treat the business of spreading innovations as a skill that improves over time.* As a beginner, start small. With experience, you can influence ever-larger communities.

### **Consortia**

Ilkka Tuomi notes that the biggest innovations of our time—Internet, Web, Linux—have been facilitated by open consortia [9]. The purpose of a consortium is to be a (virtual) meeting place where engineers can come together to

discuss and reach consensus for the basic components that enable them to cooperate, collaborate, and interoperate. In the IETF (Internet Engineering Task Force) and W3C (World Wide Web Consortium), the basic components of interest are data formats and protocols; in any of the several Linux consortia, the basic components are the software modules making up Linux. The key factor is that the consortium is able to reach agreements without a central authority and without the encumbrances of their members' organizational bureaucracies. If enough key players join, the consortium's recommendations will become de facto standards that all members and observers can use.

Viewed from the perspective of networks, a consortium is a purposefully constructed bridge between network parts. It becomes a new hub for distributing influence.

### **Distributed Responsibility**

Many vast networks have no central authority. The responsibility for taking actions, and for coping with their consequences, moves away from hubs and toward the small groups of the network. Network mechanisms, such as consortia and mediators, help groups reach agreements and resolve disputes.

One of the great challenges of our age is "network centric operations," sometimes also called distributed operations. This means

that organizations—military and commercial—operate within a distributed communication network for their communication and coordination. (The oxymoron "network centric" expresses the tension between the network's distributivity and the tradition of hierarchical command and control.) Given that operations are evolving in this way, a question on the minds of military officers and CEOs is: Can a network be commanded?

Military doctrine offers a clear answer: it's called commander's intent. The commander expresses strategic intent of an operation and leaves the details of how to accomplish it to local units. The coordination is maintained by a lot of communication between commander and local units. All units know the intent and work on it autonomously. They feed their results back to the commander, who coalesces them into new strategies and new intents. The same model applies in many large companies.

Using a military example, the early operations in Afghanistan were conducted by Special Operations Forces, which consisted of many highly decentralized local units commanded by captains and majors. The units knew the intent of their operations, but had full authority to define and execute local actions. They were able to blend with the community, get intelligence, and track down enemies. Later, the full military brought with it the centralizing

# The Profession of IT

tendencies of hierarchical command, taking authority back from the captains and majors, and slowing military progress considerably.

Perhaps the most difficult lesson for a commander or CEO is to leave the network alone after the intent is communicated and the right people chosen for the local units. The commander has to be willing to trust the local units and not try to “micromanage” them. The commanders and CEOs who have practiced this have consistently wound up with

the most successful, effective, agile, and innovative networks. ■

## REFERENCES

1. Barabasi, A.-L. *Linked: How Everything is Connected to Everything Else and What it Means for Business, Science, and Everyday Life*. Plume, New York, 2003.
2. Barabasi, A.-L. and Bonabeau, E. Scale-free networks. *Scientific American* (May 2003), 60–69.
3. Barabasi, A.-L., Albert, R., and Jeong, H. Mean field theory for scale-free random networks. *Physica A* 272, (1999), 173–187.
4. Buchanan, M. *Ubiquity: Why Catastrophes Happen*. Three Rivers, 2000.
5. Denning, P. The social life of innovation. *Commun. ACM* 47, 4 (Apr. 2004), 15–19.
6. Erdos, P., Graham, R., and Neseiril, J. *The Mathematics of Paul Erdos*. Springer-Verlag, 1996.
7. Lewis, T.G. Vulnerability analysis in critical infrastructure protection. *J. Information Warfare* 3, 2 (2004), 1–13.
8. Strogatz, S. Exploring complex networks. *Nature* 410, 8 (Mar. 2001), 268–276.
9. Tuomi, I. *Networks of Innovation*. Oxford, 2002.
10. Watts, D. *Six Degrees: The Science of a Connected Age*. Norton, 2003.

**PETER J. DENNING** (pjd@nps.edu) is the director of the Cebrowski Institute for information and innovation and superiority at the Naval Postgraduate School in Monterey, CA, and is a past president of ACM.

© 2004 ACM 0001-0782/04/1100 \$5.00

## STAY ON TOP OF ACM NEWS WITH MEMBERNET

### COMING IN FUTURE ISSUES OF MEMBERNET:

- Preview of Grace Hopper Celebration of Women in Computing
- Interviews with ACM's new President, VP and Secretary/Treasurer
- New career services for members
- E-voting: David Dill's speech at the Computers, Freedom, and Privacy conference
- Results of member opinion poll

And much more!

All online, in MemberNet: [www.acm.org/membernet](http://www.acm.org/membernet).