selves as hackers, and as this freewheeling group embraced the intense and geeky Stallman, he in turn embraced their hacker ethic. According to this unwritten credo, information and technical tricks must be freely shared and computer resources should be treated as a common good.

Soon after graduating from Harvard, Stallman became a full-time programmer in the AI lab. But by 1980, commercial software companies were beginning to appear, and these saw the AI lab as a ready pool of talent and software.

As they started hiring individuals away to work for them on proprietary products and deals were struck with MIT for exclusive rights to software developed at the AI lab, Stallman watched his hacker paradise crumble around him.

The profound sense of personal loss and anger this provoked in Stallman is probably the key motivation behind a series of actions that led to the foundation of the GNU Software Project in 1983 and to using the GPL to cast at least part of the hacker ethic in stone.

After this point, Williams' book begins to pale, largely because of the writer's extreme reluctance to criticize his subject directly (more than once, Williams confesses to feeling intimidated by Stallman). Rather than marshal evidence to back up his own judgments, Williams quotes others' critical comments, without enabling us to assess their validity. He also glosses over many events, dealing with one important period by simply noting, with little further ado, "Much has been made about the GNU project's struggles during the 1990- 1993 period."

Of course, it would be a Herculean task to analyze the often-obscure technical and personal issues that have marked Stallman's life and work since the founding of the GNU project. Still, this type of independent scholarship marks the difference between biography and hagiography.

That said, it is useful to have much of this material in one place. Stallman has influenced millions of programmers, either directly or by way of the GPL, and anybody interested in how and why software is written today should find this book fascinating.

It's a Small World After All

A tour of the theory of networks reveals some surprising links

BY STEVEN M. CHERRY
Senior Associate Editor

With breathtaking ambition, *Linked: The New Science of Networks* sets out to explain fundamental similarities across an extraordinary range of diverse and important networks.

Whether religious or social networks, the spread of infectious diseases, the web of scientific citations, or, of course, that network of networks, the Internet, we can identify all non-random networks as essentially similar. We can learn their fundamental properties. We can study one and apply its lessons to another. All that alone is remarkable.

It's almost as striking that the new science of networks can be explained to the layperson in 265 pages, as Albert-László Barabási, a professor of theoretical physics at the University of Notre Dame, does.

Theory through the ages

The first eight chapters of *Linked* entertainingly and deftly describe the basis of our current theory of networks. They begin with Leonhard Euler's accidental invention of the mathematical underpinning of our understanding of networks, known as graph theory, in 1736, and end with the latest developments, several of which are due to Barabási and his collaborators.

The story blossoms in the 1920s with the discovery of random networks, first by Paul Erdos and Alfréd Renyí. Take a collection of individuals (nodes). Start drawing lines (links) between any two nodes at random. When there are as many links as nodes, "a miracle happens," Barabási says. "[M]ost nodes will be part of a single cluster… Starting from any node, we can get to any other by navigating along the links." This holds true whether we are linking religious beliefs, proteins, documents, or electrons.

Three degrees of separation

A key question is how many hops it takes to get from one node to another. For 100 nodes, it might be as few as 1 or as many as 99. One is tempted to think that the average might be, let's say, 50. But for an astonishing variety of networks, it's not even half that. The average number of mutual-acquaintance links joining any two of the six billion or so people on earth, for example, proves to be no more than six, according to a number of studies. So despite the earth's large population, it really is a small world after all. (In fact, Barabási thinks the real number is closer to three; experiments show higher averages because there's no good way of knowing the shortest path between oneself and a complete stranger.)

It turns out that such so-called small-world networks- those where the average number of links needed to get from one node to another is much less than the number of nodes would suggest- are everywhere. In food webs (networks of which creatures eat which other creatures), species are usually only two meals from one another. The distance between two Web pages, measured in terms of the number of hyperlinks that have to be clicked to get from one page to the other, averages 19 links, the most of any network studied to date.

Networks that exhibit a small-world type of behavior aren't randomly organized, although they may start out that way. In a random network, the distribution of the number of links from one node to other nodes is a bell curve: some nodes turn out to have more links than average, some fewer, but most are quite near the average and none has very many. Small worlds, on the other hand, have a small subset of nodes with a very high number of links. These richly connected nodes are known as hubs and they provide shortcuts through the network, reducing the "distance," or number of links, that must be traversed to move between any two nodes.

Hubs might remind one of the 80/20 rule of thumb in economics, which suggests that, for example, 20 percent of taxpayers pay 80 percent of taxes, or 20 percent of a business's customers use 80 percent of the company's customer support. In rough terms, the distribution of links among nodes of a small-world network follows the 80/20 rule, since, for example, some 15 percent of pages account for 80 percent of all links in the Web. An 80/20 distribution does not follow a bell curve, but the curve of a function that decays exponentially, otherwise known as a power law. Power laws, Barabási says, are "the patent signatures of self-organization in complex systems."

What causes power law distributions of links in networks of Hollywood actors, a nation's power grid, and the wiring diagram of an IBM computer chip, to name just three that Barabási's group looked at in 1999? They found that power laws are inherent in every network with two properties: its nodes grow in number, and its links are not random. For the Web, for example, "When choosing between two pages, one with twice as many links as the other, about twice as many people link to the more connected page."

The combination of growth and preferential attachment (as this rich-get-richer phenomenon is called) characterizes power law networks like the Web, scientific citations, and Hollywood casting (where success in one film can greatly enhance the likelihood of an actor appearing in another). Barabási calls this the scale-free model of networks, in that the distribution of links between nodes, between clusters of nodes and even between clusters of clusters is the same. Whatever the scale on which you examine the network, the same power law applies.

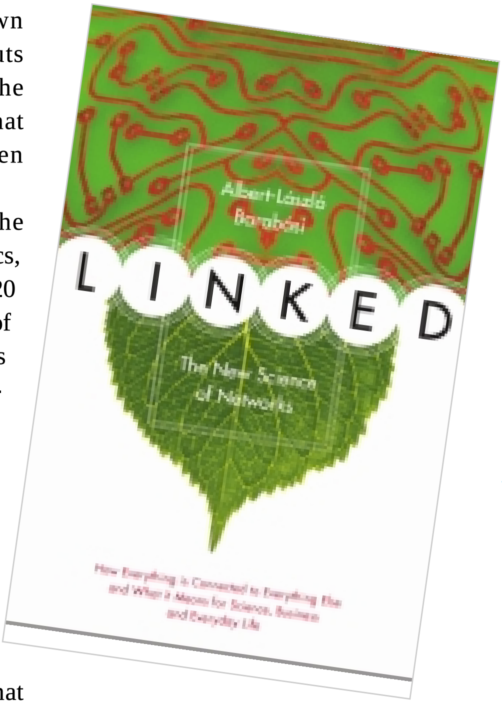

Robust, yet vulnerable

The second half of *Linked* looks in detail at some actual networks and their strengths and weaknesses, both of which, it turns out, are related to their interconnectivity. For example, networks are generally robust: they can withstand damage. Indeed, the military goal for the network that grew into the Internet was an ability to withstand attack, in the sense that surviving nodes could continue to communicate with one another by routing messages around damaged or destroyed nodes.

Ironically, the Internet is quite robust against random failure- that's a general property of scale-free networks- but not against intentional attack on its hubs. Nodes destroyed at random rarely include an important hub, and even when it is destroyed, enough other hubs and links remain for the network to work. In a theoretical experiment, Barabási's group found they "could remove as much as 80 percent of all nodes [belonging to the Internet], and the remaining 20 percent still hung together."

Conversely, a focused attack on hubs can have devastating consequences; networks like the Internet are also vulnerable to cascading failures. A network is said to be robust if, when a hub is destroyed, communication is rerouted through the remaining nodes, particularly through the remaining hubs. But those hubs can themselves fail, done in by the heavier traffic.

LINKED:
The New Science of Networks
By Albert-László Barabási
Perseus, Cambridge, Mass., 2002
ISBN 0-7382-0667-9
265 pp., US $26

Barabási looks at several such cascading failures, including those of the electrical power grid in the western United States in the summer of 1996 and the Asian economic recession of the late 1990s, which he traces to the collapse of a single Thai property development company.

Linked can only touch on these and other tantalizing examples, raising such questions as whether the spread of a disease like AIDS could be slowed by focusing limited epidemiological resources on "hubs," the presumed 20 percent of the people who engage in 80 percent of the behavior that transmits it.

Is the network of terrorism scale free? And what does that tell us about our ability to destroy it? Can the growth of cancer be stopped by targeting hub cells or molecules responsible for most of the runaway growth of a tumor? If questions like these are to be solved by networks of researchers, the publication of *Linked* dramatically increases the linkages in the web of knowledge needed to answer them. ●

Stephen Cass, *Editor*