1

ORIGINS

The "new science of networks"—an emerging field of study that we abbreviate as *network science*, is really quite old, having roots as far back as 1736. Essentially the application of mathematical graph theory to problems in a variety of fields, network science reemerged in the late 1990s as a "new science." But graph theory has been applied to practical problems since its inception in 1736, when Swiss mathematician Leonhard Euler solved the very real-world problem of how best to circumnavigate the Bridges of Königsberg, using graph theory.

Graph theorists spent the next 200 years in the backwaters of arcane mathematics. But it is difficult to keep a good idea down for long. The mathematics of graphs appeared again in the 1950s when the Hungarian and nomadic mathematician Paul Erdos (1913–1996) reestablished graph theory (and created the branch known as *discrete mathematics*) with papers on random graphs. Erdos' colorful description of mathematics as a machine for turning coffee into theorems preferred extended visits with other mathematicians to owning his own home. Today, we use the Erdos–Renyi (ER) random graph as a kind of benchmark—to compare with nonrandom graphs. The ER generative procedure for constructing a random graph marked a second historical milestone in 1959–60 (see Table 1.1).

In the late 1960s and 1970s graph theory was used by social scientists to model social networks and study the behavior of humans in groups. Stanley Milgram is credited with introducing the notion of a *small-world* network to the social science community—igniting interest in studies of how network topology might influence

Network Science: Theory and Practice. By Ted G. Lewis

TABLE 1.1 Historical Timeline of Significant Events

Date	Who	Contribution
1736	Euler	Bridges of Königsberg
1925	G. Yule	Preferential attachment, Yule–Simon distribution
1927	Kermack, McKendrick	First epidemic model
1951	Solomonoff, Rappaport	Spread of infection in random networks
1955	Simon	Power law observed in word analysis
1959	Gilbert	First generative procedure for random graph
1960	Erdos, Renyi	Random graphs
1967	Milgram	Small-world experiment
1969	Bass	Diffusion of innovation in populations—nonnetwork model
1971	Fisher, Pry	Diffusion by product substitution—nonnetwork model
1972	Bollobas	Complex graphs
1972	Bonacich	Idea of influence in social networks leading to influence diagrams
1973	Granovetter	Job-seeking networks formed clusters with "weak links" between them
1978	Pool, Kochen	First theoretical examination of small worlds
1984	Kuramoto	Synchronization of linear systems
1985	Bollobas	Publishes book on "random graphs"
1988	Waxman	First graph model of the Internet
1989	Bristor, Ryan	"Buying networks" = application of network science to model economic system
1990	Guare	Coined phrase, "six degrees of separation" = name of his Broadway play
1995	Molloy, Reed	Generation of networks with arbitrary degree sequence distribution
1996	Kretschmar, Morris	Early application of network science to spread of infectious disease = contagion driven by largest connected component
1998	Holland	Introduction of emergence in complex adaptive systems
1998	Watts, Strogatz, Faloutsos, Faloutsos	Renewed interest in Milgram's original work on small worlds, examples of clustering; first generative procedure for small world
1999	Faloutsos	Power law observed in Internet
1999	Albert, Jeong, Barabasi	Power law observed in WWW
1999	Dorogovtsev, Mendes	Small-world properties
1999	Barabasi, Albert,	Scale-free network model
1999	Dorogovtsev, Mendes, Samukhim, Krapivsky Redner	Exact solution to scale-free network degree sequence
1999	Watts	Explanation of "small-world dilemma": high clustering, low path length

(*Continued*)

TABLE 1.1 *Continued*

Date	Who	Contribution
1999	Adamic	Distance between .edu sites shown to be small-world
1999	Kleinberg, Kumar, Raghavan, Rajagopalan Tomkins	Formalized model of WWW as "Webgraph"
1999	Walsh	Difficulty of search in small worlds using local properties
2000	Marchiori, Latora,	Harmonic distance replaces path length: works for disconnected networks
2000	Broder, Kumar, Maghoul, Raghavan, Rajagopalan Stata, Tomkins, Wiener	Full Webgraph map of the WWW
2000	Kleinberg	Shows $O(n)$ search in small world using "Manhattan distance"
2000	Albert, Jeong, Barabasi	Scale-free networks are resilient if hubs are protected (Internet's "Achilles heel")
2001	Yung	Taxonomy of applications of small-world theory to: SNA, collaboration, Internet, business, life sciences
2001	Pastor-Satorras, Vespignani	Claim no epidemic threshold in scale-free networks; Internet susceptible to SIS viruses
2001	Tadic, Adamic	Use of local information can speed search on scale-free networks
2002	Levene, Fenner, Loizou, Wheeldon	Enhanced Webgraph model concluded structure of the WWW couldn't be explained by preferential attachment alone
2002	Kleinfeld,	Claims Milgram experiments not well founded: small-world social network is an "urban myth"
2002	Wang, Chen, Barahona, Pecora, Liu, Hong, Choi Kim, Jost, Joy	Sync in small worlds equivalent to stability in coupled system
2003	Wang, Chakrabarti, Wang, Faloutsos	Showed spread of epidemics determined by network's spectral radius, largest eigenvalue of connection matrix
2003	Virtanen	Complete survey of network science results up to 2003
2003	Strogatz	Synchronization of crickets, heartbeats
2005	NRC	Definition of network science
2006	Atay	Synchronization in networks with degree sequence distribution—application to networks
2007	Gabbay	Consensus in influence networks—linear and nonlinear models

human behavior—and the reverse. The "small-world dilemma" was the subject of vigorous study throughout this period. Why is it, social scientists asked, that humans are able to connect to one another through an extremely small number of intermediaries, even as the size of a population grows?

Milgram's famous "six degrees of separation" experiment suggested that the *distance* between two people selected at random from the entire population of the United States is approximately six intermediaries. In Milgram's experiment, volunteers in Kansas and Nebraska were asked to forward a letter to an unfamiliar target living in Cambridge and Boston, Massachusetts. Not knowing the target person, recipients forwarded the letter to an acquaintance closer to the target than themselves. Many of the letters were lost, but of the ones that eventually reached their target addresses, the number of hops along the chain of intermediaries ranged from 2 to 10, with an average of 5.2. Hence the notion of a small world and six degrees of separation was born.

Network science took its third, and current step toward becoming a scientific discipline of its own in the late 1990s when a number of scientists in other fields began to use networks as models of physical and biological phenomena. In particular, the pioneering work of Duncan Watts, Steven Strogatz, and Albert-Laszlo Barabasi stimulated renewed interest in mathematical analysis of networks as applied to the physical world. Watts equated the structure of very sparse networks with small diameter (small worlds) with a diverse number of phenomena such as phase transitions in materials, functionality of biological organisms, and behavior of electrical power grids. How could a simple graph model explain such diversity of real-world behaviors?

Strogatz studied the impact of network structure on complex adaptive systems in physics as well as explaining why hearts beat in a regular synchronized pattern in mammals, and why a certain species of firefly rhythmically chirps in unison without centralized control. It appeared that living organisms tend to synchronize their behavior without global knowledge. In this book, we show that a deep understanding of how and why network synchronization occurs in physical and biological systems also explains the conditions for arriving at a consensus by a group of people, how best to conduct product marketing campaigns, and how corporations rise to become a monopoly. Synchronization is a byproduct of the structure of "living networks."

Barabasi and students created another line of investigation with the invention of *scale-free networks*—nonrandom networks with hubs. In a number of studies of the structure of the Internet and WWW, Barabasi et al. discovered an emergent property of the decentralized Internet—that it had emerged without central planning into a structure consisting of a small number of extremely popular sites called *hubs*, and a large number of "unpopular" sites with few links. Instead of being random, like an ER (Erdos–Renyi) network, the Internet topology was very nonrandom. In fact, the probability that a site has k links obeys a power law, which drops off quickly for large k. Furthermore, they speculated that this was the result of a microrule called *preferential attachment*—that the probability a site will obtain a new link is directly proportional to the number of links it already has. Thus, the more links a site has, the more it gets—the so-called "rich get richer" phenomenon.

Scale-free networks are *extremely* nonrandom. This discovery set the stage for a plethora of publications in a diverse number of disciplines ranging from political science to sociology, biology, and physics. Why do so many natural phenomena obey the power law instead of the normal distribution or perhaps the exponential distribution? Once again, the stage was set for deep inquiry into the structure of organizations, organisms, and physical matter to explain the questions raised by the power law.

The current state of network science can best be described as "still evolving." In its modern form, it is approximately a decade old. Discoveries continue to be made on a monthly basis, which means that this book will soon be out of date! Therefore, the author has attempted to focus on the fundamentals—results that hopefully will endure for decades—rather than delve into interesting but distracting diversions.

The purpose of this chapter is to define the emerging discipline called *network science* and develop a historical timeline of key events leading up to the current state of the field. We survey that past 270 years leading up to the current state of the art, and end with a loose collection of "rules" for networks. We study the following in detail:

1. Network science can be defined in many ways. We loosely define it as the study of the theoretical foundations of network structure/dynamic behavior and its application to many subfields. Network science is both theory and application.
2. The history of network science is divided into three periods: (1) early pre–network period (1736–1966), when network science was really the mathematics of graphs; (2) a meso–network period (1967–1998), when network science was not yet called "the new science of networks," but in fact applications of networks were emerging from the research literature; and (3) the modern period (1998–present), when the pioneers of the current definition of network science set forth the fundamentals and showed that the fundamentals had meaning in the real world.
3. The key concepts or principles of network science are (at least) structure, dynamism, bottom–up evolution, autonomy, topology, power, stability, and emergence. Each of these are explained in detail in this chapter.
4. We give a new perspective on network science that links emergence of network synchronization to stability of linear coupled systems. This perspective integrates a number of concepts underlying applications like models of the spread of epidemics, dynamics of various forms of network emergence, and behaviors observed in disparate fields such as biology, physics, and marketing. We claim the underlying behavior of these applications is nothing more than special cases of the more general case of linear system stability; that is, the spread of infections, consensus building in groups, stability of electric power grids, and so on are applications of coupled linear system analysis. All of these seemingly disparate behaviors can be explained and analyzed using the tools of spectral analysis.

1.1 WHAT IS NETWORK SCIENCE?

The Committee on Network Science for Future Army Applications, commissioned by the Board on Army Science and Technology in cooperation with the National Research Council of the National Academies, defines *network science* in a number of ways, and at a number of levels of detail (National Research Council, 2005). Perhaps the simplest and most direct definition given by the NRC is "organized knowledge of networks based on their study using the scientific method." This definition is meant to distinguish network science from the various technologies that use it—for example, to separate the underlying science of networks from technologies such as the Internet.

But network science is not yet mature enough to be separated from its technological roots. The Committee discovered that each subfield using network science had a different working definition. Communication engineers think of networks as systems of routers and switches; sociologists think of networks as influence diagrams representing the social interactions among humans; marketing business people think of networks as populations of buyers; and the physicist thinks of networks as models of phase transition, magnetism, and so on. Biologists use the network metaphor to understand epidemics, genetics, and metabolic systems within cells, and power engineers think of electrical power grids. Network science appears to be in the eye of the beholder with different nomenclature, different vocabulary, and different methods of analysis in each field.

Perhaps it is easier to define what a *network* is, than what *network science* is. In this respect, the concept of a network is more universal, even though the terminology is not. The Committee describes a network "by its structure (e.g., nodes and links), and its behavior (what the network 'does' as a result of the interactions among the nodes and links)." It goes on to say, "a network is always a representation or model of observable reality, not that reality, itself." Networks are graphs that represent something real.

The operational definition above identifies two key ingredients of network science: (1) it is the study of the structure of a collection of nodes and links that represent something real, and (2) it is the study of the dynamic behavior of the aggregation of nodes and links. It asks, "What happens over time as a network evolves, and why does it happen?" The most significant results of network science seem to correlate form with function and structure with behavior. Currently, the behaviors of greatest interest are in physical, biological, and social systems. Nodes might be humans, molecules, genes, routers, transformers (in power grids), Web pages, or research publications. Links might be friendships, contagions, synapses, cables, Internet links, or bibliographical citations. In this sense, network science is an abstraction of reality—not the reality, itself. However, if the abstraction can explain the behavior of a real system, then network science is not only highly interesting but useful as well.

The structure portion of a network is easily modeled by graph theory. Specifically, the network itself can be defined in terms of a set, $G = \{N, L, f\}$, where N is a set of nodes, L a set of links, and f: $N \times N$ a mapping function that defines the structure of G—how nodes are connected to each other through links. The mapping function contains enough information to draw the graph on a planar piece of paper using dots as nodes and lines as links. But the set G is inadequate to define the second part of a network—its dynamic behavior.

The dynamic portion of a network is defined by a set of *microrules* governing the behavior of nodes and links. These rules are given at the *microlevel*, to distinguish them from *macrolevel* behaviors of networks. Specifically, microlevel rules dictate the behavior of links and nodes, and macrolevel rules dictate the emergence of global properties of a network. For example, *preferential attachment*—links are attracted to nodes with a lot of links, already—is a microrule, whereas the power law describing the *degree sequence distribution* of a network is a macrolevel rule. As network scientists, we are concerned mainly with understanding macrolevel properties by studying microrules—and sometimes the reverse.

A complete definition of network G must include both structural and behavioral information. For example, $G(t) = \{N(t),\ L(t),\ f(t)\}$ is a set-theoretic definition of network G with a dynamic dimension—$G(t)$ is a function of time t and the number, values, and mappings of nodes and links as they change with time. The actual behaviors of $G(t)$ are expressed algorithmically, typically in the form of a computer algorithm. In this book we use programming language *Java* to express microrules. Taken together, a compact definition of a network with its structural and behavioral elements is presented below.

Definition of Network

$$G(t) = \{N(t), L(t), f(t) : J(t)\}$$

where, $t =$ time, simulated or real

$N =$ nodes, also known as *vertices* or "actors"

$L =$ links, also known as *edges*

$f : N \times N =$ mapping function that connects nodepairs, yielding topology

$J =$ algorithm for describing behaviors of nodes and links versus time

We can now propose a rigorous definition of network science as the study of networks, including their structure and behaviors:

Definition of Network Science *Network science*, or the *science of networks*, is the study of the theoretical foundations of network structure/dynamic behavior and the application of networks to many subfields. Currently known subfields include *social network analysis* (SNA), *collaboration networks* (bibliographic citations, product marketing, online social networks), synthetic *emergent systems* (power grids, the Internet), physical science systems (phase transition, percolation theory, Ising theory), and life science systems (epidemics, metabolic processes, genetics).[1]

It should be clear from this definition that network science is essentially the science of systems. In addition, because networks often model complex systems, it is closely associated with the older field of complex adaptive systems. In fact, network science incorporates ideas from complex adaptive systems (emergence), chaos

[1]Ernst Ising (1900–1998) proposed the model of phase transition from paramagnetic to ferromagnetic state that bears his name. At some point, enough atoms align in the same direction to create a magnet in iron. This is the result of transition from a state of predominantly random polarity (atoms cancel one another) to a state of minimum energy or predominantly aligned atoms, hence producing an overall magnetic effect.

theory (synchronization), and mean-field theory (physics). Network science is a crossroad of sorts, pulling together ideas from its sister disciplines.

1.2 A BRIEF HISTORY OF NETWORK SCIENCE

Network science has been around for a long time, especially if graph theory is considered its genesis (see Table 1.1). But network science is more than graph theory, because of its dynamic aspect and application to a number of other disciplines. In general, network science has roots in graph theory, social network analysis, control theory, and more recently, the physical and biological sciences. In a sense, network science is the result of convergence of many other fields.

From a distance, it appears that network science has undergone (at least) two major transitions: from mathematical theory to applications of graphs, and from applications to a collection of generalizations about "things that are connected." Accordingly, we divide the brief history of network science into three periods: (1) early pre–network period (1736–1966), when network science was really the mathematics of graphs; (2) the meso–network period (1967–1998), when network science was not yet called "the new science of networks," but in fact applications of networks emerged from the research literature; and (3) the modern period (1998–present), when the pioneers of the current definition of network science set forth the fundamentals and showed that the fundamentals had meaning in the real world. In the modern period, advocates of network science began to demonstrate the universality of network science as they applied it to diverse fields that seemingly had no relationship to one another.

1.2.1 The Pre–Network Period (1736–1966)

The first known application of network science was Euler's treatment of the Bridges of Königsberg (Euler, 1736). It is significant because it established graph theory and showed that abstractions of reality can indeed be useful for solving problems in the real world. In one stroke, Euler defined the static structure of a physical system in terms of abstract mathematical objects called *vertices* (nodes) and *edges* (links). Logically reasoning on the abstract level, he showed that it was impossible for the citizens of Königsberg to parade through town and return without crossing one of its seven bridges at least twice.

Seven Bridges of Königsberg Problem Königsberg, Germany—aka Kaliningrad, Russia—is a city with part of its landmass on an island in the middle of the river Preger, and another part separated by a fork in the river. Seven bridges allow its citizens to get from any of its four landmasses to any other. Four bridges connect the banks of the river to the island; two bridges cross the forks of the river, and one bridge connects the island to the landmass located between the forks. The city fathers called on Leonhard Euler to tell them whether it was possible to parade throughout the entire city of five landmasses and cross each bridge only once. The solution to this problem is given in Chapter 2.

Leonhard Euler was a prodigious mathematician. It took the Swiss authorities 48 years after his death to publish his entire works. He remains the father of graph theory, today, and his legacy is the theoretical basis for the structural part of network science.

Mathematicians have added thousands of graph theory results since Euler laid the foundation more than 270 years ago. Graph theory has been extremely useful in computer science and electrical engineering, as well as a number of other applied disciplines. But as far as network science is concerned, the next major step came in 1925 when Yule first observed *preferential attachment* in evolution (Yule, 1925). Yule's work seems to have little to do with network science, but his idea would resurface in the 1990s as an explanation for the evolution of the Internet and WWW (World Wide Web). Preferential attachment explains why scale-free networks exist in natural and synthetic systems.

Preferential attachment is a simple emergent behavior observed in a number of disciplines. In the context of networks, it states that a network grows by adding nodes and links—not randomly, but by preference. Suppose that a network starts with three nodes connected by a single link between two of the three nodepairs. Now suppose that new nodes are added at regular time intervals; that is, let the network grow through a systematic process of connecting one node at a time to existing nodes. How should the new nodes be linked to existing nodes? Random attachment is one algorithm we might use: selecting an existing node at random and connecting it to the new node by adding a new link between the nodepair.

An alternative (and Darwinian) algorithm is as follows. Connect the new node to an existing node with probability proportional to the number of links already connected to the existing node. The number of links connecting a node is called its *degree*. This rule says to give preference to nodes with high degree. Thus, a new node is more likely to be connected to an existing node containing two links than to nodes with only one link. The new node will prefer being connected to the node with the higher degree.

Preferential attachment describes an *emergent process*—that is, a process that results in a network topology that is not apparent by examination of the local algorithm, or *microrule*. It is not at all obvious that the result of repeated application of preferential attachment will result in a network with a *degree sequence distribution* that follows a *power law*. This realization would come 70 years later, when A.-L. Barabasi and R. Albert showed how to create a scale-free network by repeated application of preferential attachment.

In 1927 another seemingly unrelated discovery occurred when Kermack and McKendrick published the first mathematical model of the spread of an infection in a biological population (Kermack, 1927). The *Kermack–McKendrick epidemic model* is a nonnetwork model, but it set the stage for two important innovations to come: (1) it explained the spread of a contagion along (social) links connecting (individual) nodes, and (2) it coincidentally described *new-product adoption—diffusion of technology*—and how product information spreads like an infectious disease throughout a social network. The first innovation is important because the Kermack–McKendrick epidemic model leads to the discovery of the laws of virus spreading in networks such as the Internet. In fact, we show that rate of spreading and persistence of an infection is determined completely by a network's topological structure as well as the infectiousness of the contagion. Thus some networks are more prone to epidemics than are others. Furthermore, understanding the relationship between network topology and the spread of an infection tells us how best to *stop* the spread.

The second significance of the Kermack–McKendrick model is its application to marketing of new products in the business world. The spread of information (advertising or "buzz") in a social network is much like the spread of an epidemic. What property of a network accelerates or retards this virus-like spread? In this case, merchandisers want to know how to increase infectiousness.

Solomonoff and Rappaport were the first to apply the ideas of epidemics to networks (Solomonoff, 1951). Thus the connection was made over 50 years ago (i.e., around the mid-1950s), but the relationship between network topology and infectiousness would have to wait for a more recent advance. Solomonoff and Rappaport assumed that the network was random. Today we know that random networks rarely exist in the real world—and when it comes to social networks, randomness is far from reality. Furthermore, we now know that the structure of a nonrandom network can have a dramatic impact on its function.

In fact, the idea of nonrandomness as a factor in behavior and the connection between preferential attachment and nonrandom distributions occurred to Simon in 1955 (Simon, 1955). Simon was aware of Yule's work on preferential attachment and the distribution of species among plants genera. Simon's observations confirm the validity of the power law in natural and synthetic systems: namely, that the distribution of word frequencies in documents, distribution of number of papers published by scientists, distribution of cities by population, distribution of wealth, and distribution of species among genera all obey a power law (Mitzenmacher, (2004). The evidence supporting nonrandomness in real-world phenomenon was mounting long before it was observed in networks. But the connection between nonrandomness in systems and graph theory was yet to be discovered.

By the midtwentieth century science reasoned that nature could be modeled as a random process and therefore as a random graph. What were the properties of random graphs that made them good models? Gilbert showed how to build a random graph by first constructing a complete graph and then deleting randomly selected links until reaching the desired number of links (Gilbert, 1959). But his cumbersome algorithm was quickly surpassed by the elegant and widely promoted algorithm of Erdos and Renyi (Erdos, 1960). The Erdos–Renyi (ER) algorithm is used today because of its simplicity. A network with n nodes is constructed by inserting a link between randomly selected nodepairs. The process is repeated until m links have been inserted.

The ER algorithm is not perfect—it can leave some nodes isolated, and unless it is slightly modified, it can insert duplicate and loop links into the network. But it has become the standard method of generating a random network by computer. The Gilbert and ER algorithms were the first *generative methods* of network creation. Many more methods have since been proposed, and in fact, we propose a dozen more in this book. Today, computer algorithms exist to generate a network of any prescribed topology, with any number of nodes and links; see Chapter 7.

By the very late 1960s network science did not exist, but its seeds were planted in disconnected and seemingly unrelated disciplines. It would take several more decades of scattered research in disparate disciplines before convergence to what we now know of as the science of networks.

1.2.2 The Meso–Network Period (1967–1998)

A stunning experiment performed in 1967 propelled network science from pure graph theory into scientific inquiry. The famous "six degrees of separation" experiment of Stanley Milgram seemed innocent enough, but in retrospect, it marked a turning point. Stanley Milgram invited human subjects from Kansas and Nebraska to participate in a "communications project" to "better understand social contact in American society." The experiment required them to send a folder across the country to a target person defined by the experimenters. Subjects were told to perform the following four steps (Yung, 2001):

1. Add your name to the roster at the bottom of this sheet, so that the next person who receives this letter will know whom it came from.
2. Detach one postcard, fill it out, and return it to Harvard University. No stamp is necessary. It allows us to keep track of the progress of the folder as it moves toward the target person.
3. If you know the target person on a personal basis, mail this folder directly to him/her. Do this only if you have previously met the target person and know each other on a first-name basis.
4. If you do not know the target person on a personal basis, do not try to contact him/her directly. Instead, mail this folder (postcards and all) to a personal acquaintance who is more likely than you to know the target person. You may send the folder on to a friend, relative, or acquaintance, but it must be someone you know on a first-name basis.

Subjects had 24 hours to forward their folder and received only a certificate of appreciation for their efforts. Most folders never made it. But folders that reached their target did so in far fewer steps than expected. Out of millions of people, folders passed through only a handful of intermediaries before reaching their intended destination. This winnowing of paths through a large population was evidence of the *small-world effect*, and the social network underlying Milgram's experiment is now known as a *small-world network*. The number of intermediaries averaged 5.2.

How could a stranger connect with another stranger in fewer than six steps? Milgram had to conclude that the fabric of society formed a nonrandom network. Instead of randomly bouncing from person to person, folders made a beeline to their destinations—and yet, the path they followed was not planned out ahead of time, nor was there any assurance that a chain of intermediaries existed between sender and receiver.

Assuming that each person knows 500 other people, on average, the odds of reaching the target person is approximately 1 in 200,000, and the number of intermediaries should have been many times larger than 6 hops, if paths taken by successful folders were truly random. Instead, successful folders reached their target in an average of 5.2 steps, or hops, from the originating person to the target person. The distance traveled in graph theory terms was 5 or 6 hops because five or six people handled a folder.

Milgram's experiment inspired a Broadway play and movie by John Guare called *Six Degrees of Separation: A Play* (Guare, 1990). The terminology took root, and the idea of small-world social networks blossomed, leading to purposeful creation of many other social networks. The "Kevin Bacon game" created by Brett Tjaden of the University of Virginia is one such example.[2] This network links actors that have appeared in the same movie together. The distance between Kevin Bacon and any other actor is equal to the number of hops from the node representing Kevin Bacon to any other node representing another actor.

Milgram concluded that the social world is much smaller than the "real world" because it took only 6 hops to link a pair of strangers, regardless of where they lived. He called this the *small-world problem*. Many decades later Watts and Strogatz would rejuvenate interest in small-world networks and introduce it to physicists and biologists. Their technical analysis of large sparse networks rigorously defined small worlds as networks with relatively short distances (hops) between node pairs chosen at random, even as the size of the network grows. Specifically, the diameter of a network increases as $\ln(n)$ while its size increases by $O(n)$, where n is number of nodes.

The small-world idea is related to the "weak ties" theory of Granovetter, who postulated that social networks contain both strong (direct) and weak ties (long-distance connections) that bind society together (Granovetter, 1973). In a cleverly titled paper, "The strength of weak ties," Granovetter suggests that social networks are held together by strong connections between friends and family members, as well as weaker, long-distance connections among casual acquaintances. This explains why it is possible to span a large sparse network in a small number of hops. The links of a social network are like freeways and streets—freeways have few intersections (nodes) and allow you to travel long distances without stopping. City streets, on the other hand, allow you to pinpoint a single person or house within a dense neighborhood. Freeways (weak ties) get you to a local neighborhood quickly, while streets (strong ties) zero in on a specific person or house.

White identified biases in Milgram's experiment and suggested modifications that lead to an average of seven intermediaries (White, 1970). Hunter and Shotland modeled the experiment as a Markov process to determine average distances between groups (Hunter, 1974). Pool and Kochen provided the earliest known theoretical analysis of small-world networks in 1978 (Pool, 1978). The flurry of publications stimulated by Milgram's experiment dropped off precipitously after 1978, only to be rejuvenated two decades later (Kleinfeld, 2002).

[2]This game is available at http://www.cs.virginia.edu/oracle/.

Meanwhile, pioneering work began in soon-to-be-related disciplines: Bass, Fisher, and Pry model new-product adoptions as the propagation of an infectious disease (Bass, 1969, 2001; Norton, 1987). This work extended the Kermack–McKendrick epidemic model to the new field of marketing and prepared the way for network-based product diffusion models. These models have proved to be powerful tools for the business of marketing. But do the epidemic models work when applied to social networks? We show that technological diffusion (adoption of new products) obeys the Bass and Fisher–Pry equations for a single product and random network. We also show that a monopoly arises from a random population because of preferential attachment. However, we discover that the Bass/Fisher–Pry model has limitations when modeling competition among products in a multiproduct network. These results have not been reported elsewhere.

In the social sciences, graph theory was being used to explain a number of other social interaction phenomena. Bonacich was perhaps the first social scientist to realize that influence in a social network could be mathematically represented using the connection matrix of the network (Bonacich, 1972). The nodes represent individuals, and weights on directed links represent the *degree of influence* one individual has on another.[3] If person A influences the decision of person B, and person B influences the decision of person C, and so on, what is the overall effect on group consensus? Will a chain of influences propagate through a network and eventually settle down to a common value? Bonacich claimed that consensus would eventually be reached, and proposed that the consensus be computed by raising the weighted connection matrix to the nth power, where n is the number of nodes in the network. As it turns out, influence spreading in a social network is not as simple as Bonacich's model suggests, but Bonacich initiated a line of research that continues, today.

Marketing gurus note that highly connected people are *superspreaders*—people who accelerate the spread of buzz—new-product information, simply because they are highly connected (Rosen, 2000). But social scientists have long known of the power of the middleperson or intermediary—actors who connect other actors. If the only way actor A can communicate with actor C is by going through actor B, then B has power over A and C. Thus, social scientists define *betweenness* as the number of paths that must run through an actor to connect with other actors. So, in addition to connectedness, an actor derives influence by serving as an intermediary. Does betweenness give an actor more influence than connectedness? This is the question we address in the study of influence networks. See Chapter 10.

Influence spreading—whether it is for marketing a product, spreading a disease, or achieving consensus in a group—is a form of signal propagation. The signal travels along links and affects nodes in some way. For example, the value of a node might be the average over the value of adjacent nodes. In a *Kirchhoff network*, the value of a node is equal to the difference between the sum of values from input and output links. Regardless of the local microrule for assigning values to nodes, the concept of signals flowing through networks appears to be an underlying mechanism common to epidemiology, synchronicity, influence, and consensus in groups.

[3]Weights are fractions: $0 \leq \text{weight} \leq 1$.

More rigorously, a network can be regarded as a coupled system. The system is composed of nodes that take on values called *states*, and links that establish the inputs and outputs to nodes. The state of the network is the union of the states of all of its nodes. Signals (values) propagate along links, from node to node, and alter the node's states. If we plot the change in state versus time, we might observe oscillations, dampening, or convergence to a certain state—the so-called *fixed point*—and remain there, forever. Under what conditions does a network oscillate or converge? This is a universal question, which we address in Chapters 10 and 12.

We show that spread of epidemics, synchronization of biological systems, consensus in a social network, and diffusion of new products are all different forms of synchronization in a network. A network is said to *sync* when the value of its nodes reach a *fixed point*—a value that ceases to change once reached. We answer the question "What properties or conditions are necessary and sufficient for a network to sync?" The answer leads to a general theory of stability in networks.

Kuramoto provided a mathematical basis for studying synchronization in coupled linear systems (Kuramoto, 1984). His work would influence Strogatz a decade later, and have a major impact on the convergence between network science and control theory. For example, Kuramoto's work led Strogatz to observe automatic synchronization in small-world networks. Strogatz claimed that synchronization is simply a property of all small worlds (Strogatz, 2003). This turns out to be false, but it stimulated further study into network models of various biological systems that automatically synchronize. Now we know that other conditions must exist in networks for them to sync.

By 1998, the fundamentals of network science had been established, but the explosive interest in application of the fundamentals to real-world systems was yet to come. The rapid rise of the Internet beginning in the early 1990s provided an incentive for a new generation of researchers looking to formalize the human-created, but highly decentralized, Internet phenomena. Waxman proposed a static graph theory model of the Internet in 1988 (Waxman, 1988). We call such networks *Webgraphs*, because they use graph theory to understand the World Wide Web (WWW). It would take another decade before researchers would make the connection between graph theory and the dynamic growth of the Internet and WWW. But once they did, network science came of age.

1.2.3 The Modern Period (1998–Present)

Networks have static and dynamic properties. Static graph properties such as diameter, average path length, connection matrix, and cluster coefficient provide a means for classification of the network. The *degree sequence distribution* of a network, for example, is a histogram of percentage of nodes with degree d versus d. A random network has a degree sequence distribution that obeys a binomial distribution, and a scale-free network's degree sequence obeys a power law. A small-world network has relatively small diameter and average path length, and a scale-free network has hubs. These are various ways of classifying networks on the basis of their structure. But classification according to static structure is not enough.

Networks also have dynamic properties such as fixed points when they sync, and preferences for linking nodes together when preferential attachment is operational. Dynamic networks evolve. Starting at some initial state, a dynamic network may transition to a second, third, fourth, and higher state until either cycling back or reaching a final state—its *fixed point*. The evolution from initial state to some future state is a form of *emergence*. Therefore, networks that reach a fixed point are different from networks that oscillate, forever. In this way we can classify networks according to dynamic properties.

Network science is the study of both static and dynamic properties of networks. In this book we focus on emergence as a method of understanding and characterizing the dynamic part of network science. We further divide the emergent approach into two parts: the part that alters the topology of a network (e.g., preferential attachment) and the part that alters the state of a network (e.g., synchronization). We show that a network of any desired structure can be generated as a fixed point of an emergent process whereby links are rewired until the desired topology emerges. In the final chapters we show that stability emerges out of chaos in networks with certain initial conditions and certain static properties.

Holland defines emergence in general terms as "much coming from little" (Holland, 1998). This is precisely what happens when a network evolves. A series of microlevel changes accumulate over time, until a macrolevel change in the network is realized. *Emergence* means that a major change in global properties comes from many small changes at the local level. Emergence plays a major role in the modern interpretation of network science. In a sense, it completes the puzzle of what defines the field.

Watts and Strogatz rekindled interest in small-world networks by showing the universality—and utility—of the small-world model (Watts, 1998; 1999a; 1999b). They proposed a simple emergence process—called a *generative procedure*—for constructing a small-world network. The idea is simple but brilliant: initially constructing a graph with regular structure. Next, apply rewiring to every link with probability p, such that $p*m$ links are randomized (redirected to a randomly selected node). Parameter p is the rewiring probability, and m is the number of links in the original regular graph.

The Watts–Strogatz generative procedure is tunable—increasing p also increases the randomness of the small world. This means we can produce a network with any desired level of randomness—*entropy*, if you will—by adjusting rewiring probability p—a low rewiring probability generates a nonrandom structure, and a high probability generates a random structure. The algorithm is an example of the first kind of emergence—evolving the structure, not the state, of the network.

A Watts–Strogatz network falls between a random network and nonrandom network. At a certain rewiring probability called the *crossover point* (also *length scale*), the network transitions from mostly structured to mostly random network. Typically, the crossover point is very small—on the order of $p^* = 2–3\%$. Watts and Strogatz attached a meaning to the crossover, suggesting that it corresponds to phase transition in materials (change in state from liquid to solid, magnetism, etc.). Thus, the crossover point is also known as the *phase transition threshold*.

Small worlds were no longer restricted to social networks as in the Milgram experiment after Watts and Strogatz showed how to generate an arbitrarily small world network. Moreover, small-world networks were observed in both natural and synthetic systems, which suggested that they might somehow be universal models. What does a database of film actors, the electric power grid of the western United States, and the neural network of the nematode worm *C. elegans* (*Caenorhabditis elegans*) have in common? They are all small worlds. This had to be more than a coincidence.

The small world was not the only classification of networks that seemed universal. The year 1999 was full of discoveries: M. Faloutsos, P. Faloutsos, and C. Faloutsos observed a power law in their graph models of the Internet, and Albert, Jeong, and Barabasi obtained similar results for the WWW (Faloutsos, 1999; Albert, 1999). Small worlds had characteristic short path length, but power-law networks had *hubs*—nodes with extremely high degree. Barabasi and students found that the degree sequence distribution of many synthetic and natural networks followed a power law. A network that follows a power-law distribution means that it has a hub, and many other nodes with many fewer links than the average. This lopsided preference for hubs seemed counter to nature, which typically follows a normal distribution. The trouble was that researchers were discovering too many systems structured according to the power law to ignore.

Barabasi and Albert generalized the concept of nonrandom networks with hubs and provided a generative procedure for producing *scale-free networks* (Barabasi, 1999). The name came from observing that a function $f(x)$ *scales* if $f(ax) = a'f(x)$, which is what a power law does. Therefore, if the degree sequence distribution obeys the power law, $h(x) = x^{-q}$, then it is clear that $h(ax) = (ax)^{-q} = (a^{-q})h(x) = a'h(x)$. The important contribution, however, isn't the name, but the observation that scale-free networks exhibit a sharp decline in frequency of nodes with degree d, as d increases.

Realization of the importance of small-world and scale-free networks created a feeding frenzy among mathematicians, physicists, and social scientists from 1999 through 2002. Dorogovtsev, Mendes, Samukhim, Krapivsky, and Redner derived an exact formula for the power law of a purely scale-free network and showed that it describes many biological systems (Dorogovtsev, 2000; 2002a; 2002b; 2003). Kleinberg, Kumar, Raghavan, Rajagopalan, and Tomkins suggest the term *Webgraph* to describe network models of the WWW (Kleinberg, 1999). Broder, Kumar, Maghoul, Raghavan, Rajagopalan, Stata, Tomkins, and Wiener were the first to fully map the WWW as a Webgraph and discover its structure. It isn't random. Kleinberg gives a formal explanation for Milgram's experiment based on the "Manhattan distance" between source and target nodes. The "Manhattan distance" is defined as the number of blocks, traversed along streets in Manhattan, New York, between source and destination intersections. Kleinberg showed that it takes only $O(n)$ steps to navigate such a small world (Kleinberg, 2002a).

By the end of this flurry of discovery, the basics of network science were firmly in place. But what are these models good for? If scale-free and small-world topologies are as universal as the mathematicians and physicists claimed, then it should be easy to find examples in natural and synthetic systems. Furthermore, if scale-free and small-world topologies have *profound* meaning, we should be able to derive generalizations or universal truths from the theory. In fact, this is what happened.

Albert, Jeong, and Barabasi observed that scale-free networks were extremely resilient against random attacks, but extremely vulnerable to systematic attacks on hubs (Albert, 2000). This makes sense—a random attack will most likely strike and destroy a node with only a few links, because a scale-free network has many such nodes. On the contrary, since hubs are rare, it is unlikely that a hub is attacked. But a hub has many links, and so its demise damages a large percentage of the network. Let p_c be the fraction of damaged nodes that dismantles the network. When p_c is high, the network is resilient, because many nodes must be knocked out to dismantle the network. When it is low, the network is vulnerable. In simulations, Albert et al. found threshold values of $p_c = 28\%$ for random networks versus nearly 99% for scale-free networks under random attacks; that is, a random network dismantles when an average of 28% of its nodes are damaged. But the tables are turned when hubs are systematically attacked—only 18% of the nodes need to be attacked to dismantle a scale-free network. Thus, a scale-free network is more vulnerable than a random network when its hubs are targeted.

The experiment of Albert et al. raises a question we answer in this book: What is the meaning of resiliency (risk) in a network? We extend the results of Albert et al. and assign a risk property to any arbitrary network, based on the value of nodes, their degree, and the risk formula: $R = T * V * C$, where R is risk, T is threat probability, V is vulnerability probability, and C is consequence or damage. We show that any network can be optimally protected from dismantling if target-hardening resources are deployed to nodes and links according to an algorithm proposed by Al-Mannai and Lewis (Al-Mannai, 2007). The Al-Mannai–Lewis algorithm protects high-value hubs first and lower-valued nodes last.

In related work, Pastor-Satorras and Vespignani observed that populations forming a scale-free network have no minimum epidemic threshold that prevents an infectious disease from recurring (Pastor-Satorras, 2001). Once an infection enters a network, it rises and falls repeatedly. Persistent epidemics are real—they occur in human networks as well as on the Internet. If the Internet is a scale-free network, then what is to prevent persistent viruses from infecting the Internet, permanently?

Wang and coworkers showed the initial claim of Pastor-Satorrus to be generally false (Wang, 2003a). Instead, persistence of an infection is determined not by the network's degree sequence but by its *spectral radius*, which is defined as the largest nontrivial eigenvalue of a network's connection matrix. Therefore, network topology determines its susceptibility to epidemics, but not because it is scale-free. This profound result has significant implications for fighting both kinds of viruses—Internet and human. It also has implications for the product marketer.

The network model of systems is proving to have a profound impact on understanding resiliency, risk, epidemics, and social interactions among people in groups. It is also proving to be revolutionary in understanding chaotic behavior of coupled linear systems. Wang, Chen, Barahona, Pecora, Liu, Hong, Choi, Jost, Joy, and others applied linear systems analysis to arbitrary networks and showed that the stability of any network is a function of its topology (Wang, 2002a, 2002b, 2002c; Barahona, 2002; Liu, 2002, 2003, 2004a, 2004b; Hong, 2002; Jost, 2002). We extend these results to several classes of networks—the Atay network, a kind of dynamic network studied by Atay, which uses a local averaging algorithm to compute the state of nodes (Atay, 2006), and a new class of networks called *Kirchhoff networks*, which attempt to stabilize the value of its nodes by maintaining Kirchhoff's first law.

Emergence of fixed-point solutions (synchronization) in a dynamic network has become a powerful and general tool for understanding a number of natural phenomena. Strogatz claimed that the beating of a human's heart, the chirping of crickets, and other biological systems naturally sync because they are small-worlds (Strogatz, 2003). But we show that network synchronization has little to do with small-world topology, and everything to do with the Laplacian of the connection matrix, and the length of circuits within the network. We examine stability of Atay networks, which have applications in marketing, as well as understanding the chirping of crickets, and show that sync is a property of emergent networks that contain triangle-shaped subgraphs. As it turns out, small-world networks have an abundance of triangular subgraphs! Moreover, a network can be synchronized, by adding a triangular subgraph to one of its nodes.

The picture is more complex for *Kirchhoff networks*—defined as directed-flow networks where the state of every node is the sum of the values of incoming links minus the sum of values of outgoing links. Kirchhoff networks are abstractions of electrical power grids, electromechanical feedback systems, air traffic routes, and so on. Hence, it is important that they not blow up—but instead, stabilize quickly after a disruption or sudden change in the value of one or more nodes. Curiously, Kirchhoff networks exhibit wildly chaotic behavior followed by sudden synchronization when the lengths of circuits within the network are relatively prime numbers. This has important applications to the design of self-stabilizing power grids, self-correcting transportations systems, and perhaps self-repairing biological systems.

A similar question of stability leading to synchronization—or not—can be observed in social networks. Consider a group of people attempting to arrive at a consensus (Gabbay, 2007). Each member of the group starts with an initial position, and attempts to influence his or her neighbor's position by exerting a positive or negative influence on nearest neighbors. Influence spreads like an epidemic, but rather than infecting or not, the influence sways the position of adjacent nodes. For example, each node of the social network might be adjusted to equal the average value of its neighbors.

Propagation of influence in a network is much like propagation of a signal in a coupled linear system. If the network is a kind of Atay network, nodes take on a value in the interval $[-1, +1]$, representing disagreement or agreement, or some fraction in between. As influence spreads, each node value changes. If all nodes reach the same value, after some period of time, we say that the network has reached a consensus. If nodes never reach a consensus, we say that the network diverges.

The question we address is, "Under what conditions will an influence network arrive at a consensus?" This problem is related to the more general problem of under what conditions will an arbitrary network synchronize. We show that consensus is reached when the influence network's largest nontrivial state matrix eigenvalue is bounded by one, and there are no conflicts in the network. Curiously, the most influential node in the network is the one less influenced by other nodes.

Network science has achieved more recent results in the field of marketing, and understanding competition among corporations. We show that preferential attachment in a simple random network obeys the Bass and Fisher–Pry equations, but technological diffusion is more complex in multiproduct, multicompetitor networks. Simple models of preferential attachment, value proposition, early-stage or late-stage market, and combinations of these competition models lead to monopolies in general, and oligopolies under certain conditions. Specifically, we show that it is possible for niche players to coexist with a monopoly, under most assumptions. This is an area of research still ripe for more investigation.

For more details and in-depth analysis of historical events in the brief history of network science, the reader is advised to study several excellent surveys of network science. Virtanen surveys the complete field prior to 2003 (Virtanen, 2003). The thesis by Voon Joe Yung gives a nonmathematical introduction and includes a copy of the Milgram experiment folder (Yung, 2001). The numerous papers by Mark Newman and Duncan Watts provide compact tutorials on the fundamentals of small worlds (Watts, 1999a; Newman, 2000b).

1.3 GENERAL PRINCIPLES

The "modern period" is defined by the convergence—commencing in the late 1990s—of several complementary and interrelated fields—with more yet to come. Graph theory, social network analysis, epidemic modeling, market competition modeling, and synchronization of physical and biological systems are all different aspects of network science. In fact, it may be too early to make generalizations or identify principles of network science. At some risk of being shortsighted, the author makes the following observations:

Characteristics of Network Science in the Modern Period

1. *Structure*. Networks have structure—they are not random collections of nodes and links. For example, the structures of electrical power grids, online social networks, and the nervous system of the *C. elegans* nematodes are not random, but instead have a distinct format or topology. This suggests that function follows form—many real-world phenomena behave the way they do because of their network structure.
2. *Emergence*. A network property is emergent if it changes by a factor of 10 as a consequence of a dynamic network achieving stability. In other words, emergence is a network synchronization issue—stable networks transition from one state to another until they reach a fixed point, and stay there. The fixed point is a new configuration for the network with corresponding order-of-magnitude

change in a certain property. For example, on the Internet, a group of teenagers will form a social clique 10 times larger than expected from purely random behavior, because the clique is formed by preferential attachment; the top 1% of Americans make 10 times more money than the average American; the popularity of a few movie stars is 10 times greater than the popularity of the average movie star; and the largest cities in the world are relatively rare, and are 10 times larger than the average city. In each of these examples, the "hub property" or concentration is a consequence of some instability working its way through the network as the network achieves a new fixed point. This is the impetus behind online social networks that begin with nothing, and end up with millions of subscribers. However, it is not always clear what ingredients go into online social networking to cause explosive growth. Likewise, it is not always obvious what motivation causes an order-of-magnitude change in a network's property.

3. *Dynamism*. Network science is concerned with both structure and dynamic behavior of networks. Dynamic behavior is often the result of emergence or a series of small evolutionary steps leading to a fixed-point final state of the system. The Internet, many biological systems, some physical systems, and most social systems are growing and changing networks. One must understand their dynamic properties in order to fully understand these systems. Analysis of only their static structure, such as degree sequence, is not sufficient to understand the network. For example, network synchronization, such as in the case of chirping crickets, is a consequence of the dynamism of each cricket, as well as the structure (triangular subgraphs) of the social network of crickets.
4. *Autonomy*. A network forms by the autonomous and spontaneous action of independent nodes that "volunteer" to come together (link), rather than through central control or central planning. Structure and function arise out of chaos, more as a result of serendipity than determinism. Examples are formation of large conglomerates from the merger of small companies; emergence of large cities from small communities; and formation of global telecommunication systems from linking of many smaller, local, independent operators. The initial configuration of a network may be premeditated, but over time, the network either "decays" with the onset of some form of entropy, or adapts and changes via the absorption of energy. For example, a highway system will either decay and fall into disrepair, or improve and grow through the expenditure of effort to repair, extend, increase its capacity, and so on.
5. *Bottom–Up Evolution*. Networks grow from the bottom or local level up to the top or global level. They are not designed and implemented from the top down. This can also be regarded as a form of distributed control where network evolution is a consequence of local rules being applied locally without any centralized control. Even if the initial structure of a network is the result of a premeditated design, networks evolve and change as a consequence of their dynamism. Examples of "unplanned systems" are formation of the Internet from local networks, formation of the electrical power grid from local utilities, and formation of highway systems from local roads and animal trails.

6. *Topology*. The architecture or topology of a network is a property that emerges over time as a consequence of distributed—and often subtle—forces or autonomous behaviors of its nodes. A network is dynamic if its topology or other properties change as a function of time. Thus, topology (structure) is a consequence of Darwinian forces that shape the network. For example, scale-free networks (networks with dominant hubs) emerge from the force of "preferential attachment" (economics), unintended consequences (regulatory law), such as the vulnerability of electrical power grids as a consequence of government deregulation, or "hidden order" of decentralized infrastructures emerging from complex adaptive systems such as the rise of the Internet, formation of metropolitan civilizations, or creation of monopolies like the Microsoft Corporation.
7. *Power*. The power of a node is proportional to its degree (number of links connecting it to the network) influence (link values); and *betweenness* or *closeness*; the power of a network is proportional to the number and strength of its nodes and links. For example, Metcalf's law states that the power of a network is proportional to the square of the number of nodes it contains [e.g., the maximum number of links that a network with n nodes can contain is $n(n-1)/2$, which is approximately n^2]. The influence a person exerts on a group is proportional to the position, number, and power of colleagues the person has within the group, such as the person's connectivity. The power of a corporation, within an industry or market, is proportional to the number of customers (links) that it has, or its intermediary position within the industry. Power is a subtle but important organizing principle in most networks, but it is often called something else, such as influence, signal strength, or infection rate.
8. *Stability*. A dynamic network is stable if the rate of change in the state of its nodes/links or its topology either diminishes as time passes or is bounded by dampened oscillations within finite limits. For example, the regular and rhythmic beating of an animal's heart is controlled by a stable network of nerves that regulate the pacemaker, the loss of a power plant in the electrical power grid stabilizes quickly by switching from one source to another without disruption in supply, or the loss of a coworker causes short-term reallocation of responsibility without organizational failure.

This book approaches the subject of network science from the topology perspective, first, and then the dynamic or emergent viewpoint, second. Does topology follow function, or the reverse? Does a dynamic network behave and function in a certain way because of its topology, or does it derive its topology from its function? The following chapters provide a lot of evidence to support the "form follows function" perspective. Sparse small-world networks appear to model human social networks, because humans have limited capacity to know a large number of other humans. Scale-free networks appear to model economic constructs such as the Internet and monopolies within industrial segments, because preferential attachment is essentially the law of increasing returns of economics. Networks tend to be structured in such a way that synchronization is more likely than chaos, simply because unstable systems cannot survive in nature.

It seems logical, then, to approach the field of network science from the ground up. First, we review the basics of graphs. This will require definitions and terminology (see Chapters 2 and 3). Then, the generative procedures for producing random, small-world, scale-free, and arbitrary networks of any topology are provided and studied in Chapters 4–7. Chapter 8 extends epidemiology to the network, Chapter 9 begins the development of a unified theory of network stability, Chapter 10 applies this theory to social network analysis and group consensus, and Chapters 11–13 explore applications in greater detail. Many more chapters remain to be written, hopefully by the reader!