Chapter One
Introduction

Networks are everywhere. From the Internet and its close cousin the World Wide Web to networks in economics, networks of disease transmission, and even terrorist networks, the imagery of the network pervades modern culture.

What exactly do we mean by a network? What different kinds of networks are there? And how does their presence affect the way that events play out? In the past few years, a diverse group of scientists, including mathematicians, physicists, computer scientists, sociologists, and biologists, have been actively pursuing these questions and building in the process the new research field of network theory, or the "science of networks" (Barabási 2002; Buchanan 2002; Watts 2003).

Although it is still in a period of rapid development and papers are appearing daily, a significant literature has already accumulated in this new field, and it therefore seems appropriate to summarize it in a way that is accessible to researchers unfamiliar with the topic. That is the purpose of this book. We begin by sketching in this introductory chapter a brief history of the study of networks, whose beginnings lie in mathematics and more recently sociology. We then place the "new" science of networks in context by describing a number of features that distinguish it from what has gone before, and explain why these features are important. At the end of the chapter we give a short outline of the remainder of the book.

1.1 A BRIEF HISTORY OF THE STUDY OF NETWORKS

The study of networks has had a long history in mathematics and the sciences. In 1736, the great mathematician Leonard Euler became interested in a mathematical riddle called the Königsberg Bridge Problem. The city of Königsberg was built on the banks of the Pregel River in what was then Prussia,[1] and on two islands that lie in midstream. Seven bridges connected the land masses, as shown in Figure 1.1. (There are many more than that today.) A popular brain-teaser of the time asked, "Does there exist any single path that crosses all seven bridges exactly once each?" Legend has it that the people of Königsberg spent many fruitless hours trying to find such a path before Euler proved the impossibility of its existence. The proof, which perhaps seems rather trivial to us now, but which apparently wasn't obvious in 1736, makes use of a *graph*—a mathematical object consisting of points, also called *vertices* or *nodes*, and lines, also called *edges* or *links*, which abstracts away

---

[1]Today Königsberg lies in Russia and is called Kaliningrad.
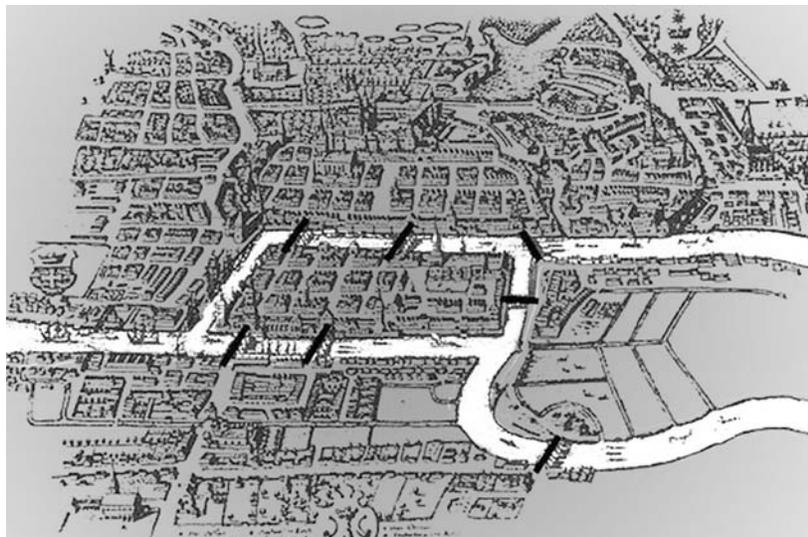
FIGURE 1.1   A map of eighteenth century Königsberg, with its seven bridges highlighted.

all the details of the original problem except for its connectivity. In this graph there are four vertices representing the four land masses and seven edges joining them in the pattern of the Königsberg bridges (Figure 1.2). Then the bridge problem can be rephrased in mathematical language as the question of whether there exists any *Eulerian path* on the network. An Eulerian path is precisely a path that traverses each edge exactly once. Euler proved that there is not, by observing that, since any such path must both enter and leave every vertex it passes through, except the first and last, there can at most be two vertices in the network with an odd number of edges attached. In the language of graph theory, we say that there can at most be two vertices with odd *degree*, the degree of a vertex being the number of edges attached to it.[2] Since all four vertices in the Königsberg graph have odd degree, the bridge problem necessarily has no solution. The problem of the existence of Eulerian paths on networks, as well as the related problem of *Hamiltonian paths* (paths that visit each vertex exactly once), is still of great interest to mathematicians, with new results being discovered all the time.

Many consider Euler's proof to be the first theorem in the now highly developed field of discrete mathematics known as *graph theory*, which in the past three centuries has become the principal mathematical language for describing the properties of networks (Harary 1995; West 1996). In its simplest form, a network is nothing more than a set of discrete elements (the vertices), and a set of connections (the edges) that link the elements, typically in a pairwise fashion. The elements

---

[2]Within physics some authors have referred to this quantity as the "connectivity" of a vertex, and the reader will see this usage in some of the papers reproduced in this book. As the word connectivity already has another meaning in graph theory, however, this choice of nomenclature has given rise to some confusion. To avoid such confusion, we will stick to standard usage in this book and refer to the degree of a vertex.
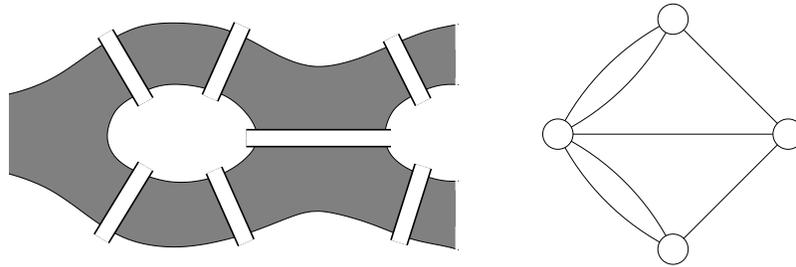
FIGURE 1.2   Left: a simplified depiction of the pattern of the rivers and bridges in the Königsberg bridge problem. Right: the corresponding network of vertices and edges.

and their connections can be almost anything—people and friendships (Rapoport and Horvath 1961), computers and communication lines (**Faloutsos *et al.* 1999**), chemicals and reactions (**Jeong *et al.* 2000; Wagner and Fell 2001**), scientific papers and citations (**Price 1965**; Redner 1998)[3]—causing some to wonder how so broad a definition could generate anything of substantive interest. But its breadth is precisely why graph theory is so powerful. By abstracting away the details of a problem, graph theory is capable of describing the important topological features with a clarity that would be impossible were all the details retained. As a consequence, graph theory has spread well beyond its original domain of pure mathematics, especially in the past few decades, to applications in engineering (Ahuja *et al.* 1993), operations research (Nagurney 1993), and computer science (Lynch 1996). Nowhere, however, has graph theory found a more welcome home than in sociology.

Starting in the 1950s, in response to a growing interest in quantitative methods in sociology and anthropology, the mathematical language of graph theory was coopted by social scientists to help understand data from ethnographic studies (Wasserman and Faust 1994; Degenne and Forsé 1999; Scott 2000). Much of the terminology of social network analysis—actor centrality, path lengths, cliques, connected components, and so forth—was either borrowed directly from graph theory or else adapted from it, to address questions of status, influence, cohesiveness, social roles, and identities in social networks. Thus, in addition to its role as a language for describing abstract models, graph theory became a practical tool for the analysis of empirical data. Also starting in the 1950s, mathematicians began to think of graphs as the medium through which various modes of influence—information and disease in particular—could propagate (**Solomonoff and Rapoport 1951**; **Erdős and Rényi 1960**). Thus the structural properties of networks, especially their connectedness, became linked with behavioral characteristics like the expected size of an epidemic or the possibility of global information transmission. Associated

---

[3]Throughout this book citations highlighted in **bold text** refer to papers that are reproduced within this book.

with this trend was the notion that graphs are properly regarded as stochastic objects (**Erdős and Rényi 1960**; Rapoport 1963), rather than purely deterministic ones, and therefore that graph properties can be thought of in terms of probability distributions—an approach that has been developed a great deal in recent years.

1.2 THE "NEW" SCIENCE OF NETWORKS

So what is there to add? If graph theory is such a powerful and general language and if so much beautiful and elegant work has already been done, what room is there for a new science of networks? We argue that the science of networks that has been taking shape over the last few years is distinguished from preceding work on networks in three important ways: (1) by focusing on the properties of real-world networks, it is concerned with empirical as well as theoretical questions; (2) it frequently takes the view that networks are not static, but evolve in time according to various dynamical rules; and (3) it aims, ultimately at least, to understand networks not just as topological objects, but also as the framework upon which distributed dynamical systems are built. As we will see in Chapter 3, elements of all these themes predate the recent explosion of interest in networks, but their synthesis into a coherent research agenda is new.

Modeling real-world networks

The first difference between the old science of networks and the new is that, social network analysis aside, traditional theories of networks have not been much concerned with the structure of naturally occurring networks. Much of graph theory qualifies as pure mathematics, and as such is concerned principally with the combinatorial properties of artificial constructs. Pure graph theory is elegant and deep, but it is not especially relevant to networks arising in the real world. Applied graph theory, as its name suggests, is more concerned with real-world network problems, but its approach is oriented toward design and engineering. By contrast, the recent work that is the topic of this book is focused on networks as they arise naturally, evolving in a manner that is typically unplanned and decentralized. Social networks and biological networks are naturally occurring networks of this kind, as are networks of information like citation networks and the World Wide Web. But the category is even broader, including networks—like transportation networks, power grids, and the physical Internet—-that are intended to serve a single, coordinated purpose (transportation, power delivery, communications), but which are built over long periods of time by many independent agents and authorities. Social network analysis, for its part, is strongly empirical, but tends to be descriptive rather than constructive in nature. With the possible exception of certain types of random graph models (Holland and Leinhardt 1981; Strauss 1986; Anderson *et al.* 1999), network analysis in the social sciences has largely avoided modeling, preferring simply to describe the properties of networks as observed in collected data.

In contrast to traditional graph theory on the one hand, and social network analysis on the other, the work described in this book takes a view that is both theoretical and empirical. In order to develop new graph-theoretic models that can account for the structural features of real-world networks, we must first be able to say what those features are and hence empirical data are essential. But adequate theoretical models are equally essential if the significance of any particular empirical finding is to be correctly understood. Just as in traditional science, where theory and experiment continually stimulate one another, the science of networks is being built on the twin foundations of empirical observation and modeling.

That such an obvious requirement for scientific validity should have made its first appearance in the field so recently seems surprising at first, but is understandable given the historical difficulty of obtaining high quality, large-scale network data. For most of the past fifty years, the collection of network data has been confined to the field of social network analysis, in which data have to be collected through survey instruments that not only are onerous to administer, but also suffer from the inaccurate or subjective responses of subjects. People, it turns out, are not good at remembering who their friends are, and the definition of a "friend" is often quite ambiguous in the first place.

For example, the General Social Survey[4] requests respondents to name up to six individuals with whom they discuss "important matters." The assumption is that people discuss matters that are important to them with people who are important to them, and hence that questions of this kind—so-called "name generators"—are a reliable means of identifying strong social ties. However, a recent study by Bearman and Parigi (2004) shows that when people are asked about the so-called "important matters" they are discussing, they respond with just about every topic imaginable, including many that most of us wouldn't consider important at all. Even worse, some topics are discussed with family members, some with close friends, some with coworkers, and others with complete strangers. Thus, very little can be inferred about the network ties of respondents simply by looking at the names generated by the questions in the General Social Survey. Bearman and Parigi also find that some 20% of respondents name no one at all. One might assume that these individuals are "social isolates"—people with no one to talk to—yet nearly 40% of these isolates are married! It is possible that these findings reveal significant patterns of behavior in contemporary social life—perhaps many people, even married people, really do not have anyone to talk to, or anything important to talk about. But apparently the respondent data are so contaminated by diverse interpretations of the survey instrument, along with variable recollection and even laziness, that any inferences about the corresponding social network must be regarded with skepticism.

The example of the General Social Survey is instructive because it typifies the uncertainties associated with traditional, survey-based collection of network data. If people have difficulty identifying even their closest confidants, how can one expect to extract reliable information concerning more subtle relations? And if, in response to this obstacle, survey instruments become more elaborate and spe-

---

[4]See http://www.norc.uchicago.edu/projects/gensoc.asp.

cific, then as the size of the surveyed population increases, the work required of the researcher to analyze and understand the resulting volume of raw data becomes prohibitive. A better approach would be to record the activities and interactions of subjects directly, thus avoiding recall problems and allowing us to apply consistent criteria to define relationships. In the absence of accurate recording technologies, however, such direct observation methods are even more onerous than the administration of surveys.

Because of the effort involved in compiling them, social network datasets rarely document populations of more than a hundred people and almost never more than a thousand. And although other kinds of (nonsocial) networks have not suffered from the same difficulties, empirical examples prior to the last decade have been few—probably because other network-oriented disciplines have lacked the empirical focus of sociology. The lack of high quality, large-scale network data has, in turn, delayed the development of the kind of statistical models with which much of the work in this book is concerned. Such models, as we will see, can be very successful and informative when applied to large networks, but tend to break down, or simply don't address the right questions, when applied to small ones. As an example, networks of contacts between terrorists have been studied recently by, for instance, Krebs (2002), but they are poor candidates for statistical modeling because the questions of interest in these networks are not statistical in nature, focusing more on the roles of individuals and small groups within the network as a whole. The traditional tools of social network analysis—centrality indices, structural measures, and measures of social capital—are more useful in such cases.

Recent years, however, have witnessed a dramatic increase in the availability of network datasets that comprise many thousands and sometimes even millions of vertices—a consequence of the widespread availability of electronic databases and, even more important, the Internet. Not only has the Internet focused popular and scientific attention alike on the topic of networks and networked systems, but it has led to data collection methods for social and other networks that avoid many of the difficulties of traditional sociometry. Networks of scientific collaborations, for example, can now be recorded in real time through electronic databases like Medline and the Science Citation Index (**Newman 2001a**; Barabási *et al.* 2002), and even more promising sources of network data, such as email logs (Ebel *et al.* 2002; Guimerà *et al.* 2003; Tyler *et al.* 2003) and instant messaging services (Smith 2002; Holme *et al.* 2004), await further exploration. Being far larger than the datasets of traditional social network analysis, these networks are more amenable to the kinds of statistical techniques with which physicists and mathematicians are familiar. As the papers in Chapter 3 of this volume demonstrate, real networks, from citation networks and the World Wide Web to networks of biochemical reactions, display properties—like local clustering and skewed degree distributions—that were not anticipated by the idealized models of graph theory, and that have forced the development of new modeling approaches, some of which are introduced in Chapter 4.

Networks as evolving structures

A second distinguishing feature of the work described in this book is that, whereas in the past both graph theory and social network analysis have tended to treat networks as static structures, recent work has recognized that networks evolve over time (**Barabási and Albert 1999**; Watts 1999). Many networks are the product of dynamical processes that add or remove vertices or edges. For instance, a social network of friendships changes as individuals make and break ties with others. An individual with many acquaintances might, by virtue of being better connected or better known, be more likely to make new friends than someone else who is less well connected. Or individuals seeking friends might be more likely to meet people with whom they share a common acquaintance. The ties people make affect the form of the network, and the form of the network affects the ties people make. Social network structure therefore evolves in a historically dependent manner, in which the role of the participants and the patterns of behavior they follow cannot be ignored.

Similar statements apply to other kinds of networks as well: processes operating at the local level both constrain and are constrained by the network structure. A principal objective of the new science of networks (as dealt with by a number of papers in Chapter 4), is an understanding of how structure at the global scale (say, the connectivity of the network as a whole) depends on dynamical processes that operate at the local scale (for example, rules governing the appearance and connections of new vertices).

Networks as dynamical systems

The final feature that distinguishes the research described in this book from previous work is that traditional approaches to networks have tended to overlook or oversimplify the relationship between the structural properties of a networked system and its behavior. A lot of the recent work on networks, by contrast, takes a dynamical systems view according to which the vertices of a graph represent discrete dynamical entities, with their own rules of behavior, and the edges represent couplings between the entities. Thus a network of interacting individuals, or a computer network in which a virus is spreading, not only has topological properties, but has dynamical properties as well. Interacting individuals, for instance, might affect one another's opinions in reaching some collective decision (voting in a general election, for example), while an outbreak of a computer virus may or may not become an epidemic depending on the patterns of connections between machines. Which outcomes occur, how frequently they occur, and with what consequences, are all questions that can only be resolved by thinking jointly about structure and dynamics, and the relationship between the two.

Questions of this nature are not easily tackled, however; dynamical problems lie at the forefront of network research, where there are many unanswered questions. One class of problems on which some progress has been made, and

which is addressed in Section 5.1, is that of contagion dynamics. Whether we are interested in the spread of a disease or the diffusion of a technological innovation, it is frequently the case that contagion occurs over a network. Not only physical but also social contacts can significantly influence the probability that a particular disease or piece of information will be transmitted, and also what effect it will have. In traditional mathematical epidemiology, as well as research on the diffusion of information, it is usually assumed that all members of the population have equal likelihood of interaction with all others. Clearly this assumption requires modification once we take network structure into account. As the papers in Section 5.1 demonstrate, the particular structure of the network through which a contagious agent is transmitted can have a dramatic impact on outcomes at the level of entire populations.

1.3 OVERVIEW OF THE VOLUME

Mirroring the themes introduced above, this volume is divided into a number of parts, each of which is preceded by an introduction that outlines the general theme and summarizes the contributions of the papers in that part. Chapter 2 sets the stage by presenting a selection of papers that we feel are important historical antecedents to contemporary research. Although recent work on networks takes a distinctly different approach from traditional network studies, a careful reading of Chapter 2 reveals that many of the basic themes were anticipated by mathematicians and social scientists years or even decades earlier. Given their age, some of these contributions seem remarkably familiar and modern, occasionally to the extent that recent papers almost exactly replicate previous results. Power-law distributions, random networks with local clustering, the notion of long-range shortcuts, and the small-world phenomenon were all explored and analyzed well before the new science of networks reconstituted the same ideas in the language of mathematical physics.

Chapter 3 emphasizes the empirical side of the new science of networks, and Chapter 4 presents some of the foundational modeling ideas that have generated a great deal of subsequent interest and activity. By exploring some tentative applications of the ideas introduced in Chapters 3 and 4, Chapter 5 takes the reader to the cutting edge of network science, the relationship between network structure and system dynamics. From disease spreading and network robustness to search algorithms, Chapter 5 is a potpourri of topics at this poorly understood but rapidly expanding frontier. Finally, Chapter 6 provides a short discussion of what we see as some of the most interesting directions for future research. We hope the reader will be encouraged to strike out from where the papers in this volume leave off, adding his or her own ideas and results to this exciting and fast-developing field.