CHAPTER 1

Overview

The past decade has seen a growing public fascination with the complex "connectedness" of modern society. At the heart of this fascination is the idea of a *network* – a pattern of interconnections among a set of things – and one finds networks appearing in discussion and commentary on an enormous range of topics. The diversity of contexts in which networks are invoked is in fact so vast that it's worth deferring precise definitions for a moment while we first recount a few of the more salient examples.

To begin with, the social networks we inhabit – the collections of social ties among friends – have grown steadily in complexity over the course of human history, due to technological advances facilitating distant travel, global communication, and digital interaction. The past half-century has seen these social networks depart even more radically from their geographic underpinnings – an effect that has weakened the traditionally local nature of such structures but enriched them in other dimensions.

The information we consume has a similarly networked structure: these structures too have grown in complexity, as a landscape with a few purveyors of high-quality information (publishers, news organizations, the academy) has become crowded with an array of information sources of wildly varying perspectives, reliabilities, and motivating intentions. Understanding any one piece of information in this environment depends on understanding the way it is endorsed by and refers to other pieces of information within a large network of links.

Our technological and economic systems have also become dependent on networks of enormous complexity. This has made the behavior of these systems increasingly difficult to reason about and increasingly risky to tinker with. It has made them susceptible to disruptions that spread through the underlying network structures, sometimes turning localized breakdowns into cascading failures or financial crises.

The imagery of networks has made its way into many other lines of discussion as well: Global manufacturing operations now have networks of suppliers, Web sites have networks of users, and media companies have networks of advertisers. In such formulations, the emphasis is often less on the structure of the network itself than on its complexity as a large, diffuse population that reacts in unexpected ways to the actions of central authorities. The terminology of international conflict has come to reflect this
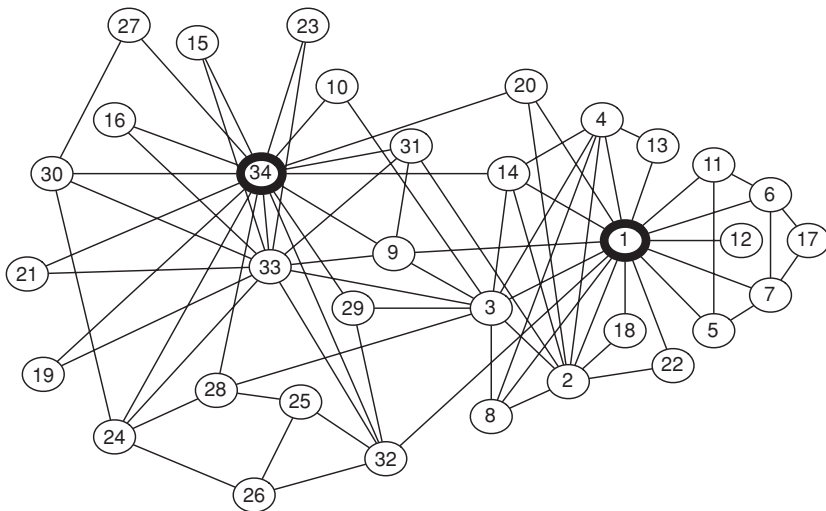
1

Figure 1.1. The social network of friendships within a 34-person karate club [421]. (Drawing from the *Journal of Anthropological Research*.)

as well: for example, the picture of two opposing, state-supported armies gradually morphs, in U.S. presidential speeches, into images of a nation facing "a broad and adaptive terrorist network" [296] or "at war against a far-reaching network of violence and hatred" [328].

1.1 Aspects of Networks

How should we think about networks, at a more precise level, to bring all these issues together? In the most basic sense, a network is any collection of objects in which some pairs of these objects are connected by *links*. This definition is very flexible: depending on the setting, many different forms of relationships or connections can be used to define links.

Because of this flexibility, it is easy to find networks in many domains, including the ones we've just been discussing. As a first example of what a network looks like, Figure 1.1 depicts the social network among thirty-four people in a university karate club studied by the anthropologist Wayne Zachary in the 1970s. The people are represented by small circles, and lines join the pairs of people who are friends outside the context of the club. This is the typical way in which networks will be drawn in this book, with lines joining the pairs of objects that are connected by links.

Later in this chapter we'll discuss some of the things one can learn from a network such as the one in Figure 1.1, as well as from larger examples such as the ones shown in Figures 1.2–1.4. These larger examples depict e-mail exchanges among employees of a company (Figure 1.2); loans among financial institutions (Figure 1.3); and links among blogs on the Web (Figure 1.4). In each case, links indicate the pairs that are connected (specifically, people connected by e-mail exchange, financial institutions by a borrower–lender relationship, and blogs via a link on the Web from one to the other, respectively).
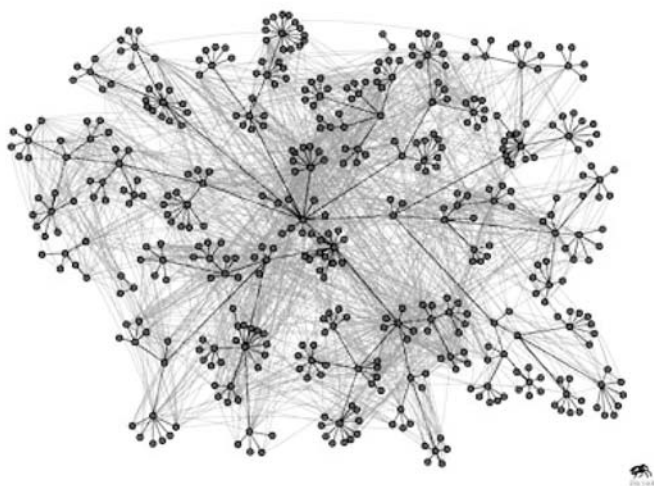
Figure 1.2. Social networks based on communication and interaction can be constructed from the traces left by online data. In this case, the pattern of e-mail communication among 436 employees of the Hewlett Packard Research Lab is superimposed on the official organizational hierarchy [6]. (Image from http://www-personal.umich.edu/ladamic/img/hplabsemailhierarchy.jpg, courtesy of Elsevier Science and Technology Journals.)
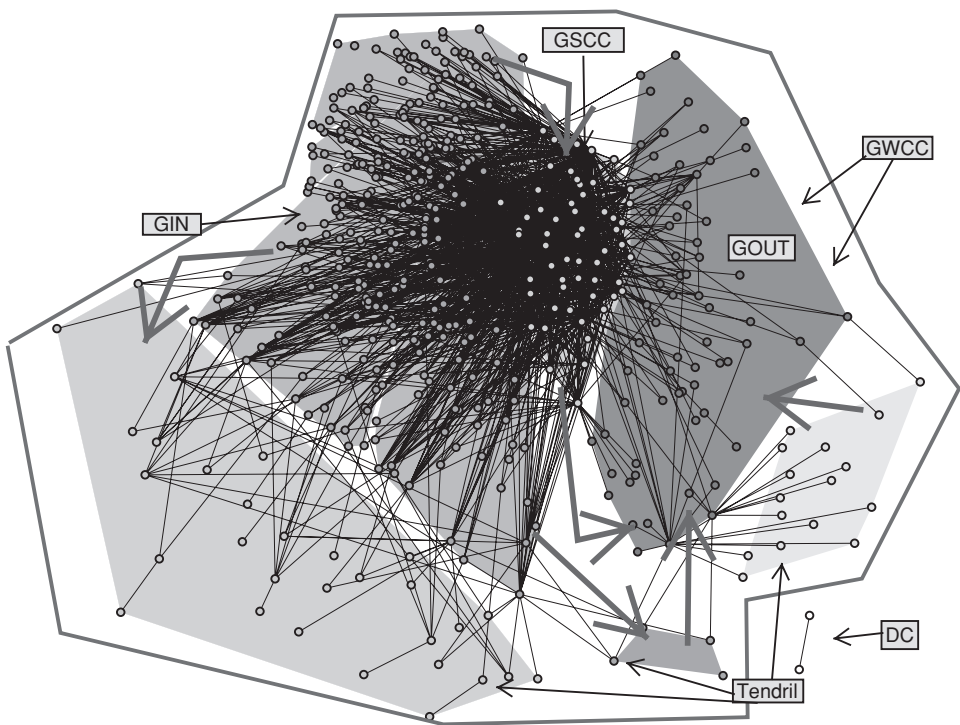

Figure 1.3. The network of loans among financial institutions can be used to analyze the roles that different participants play in the financial system and how the interactions among these roles affect the health of individual participants and the system as a whole. The network is annotated in a way that reveals its dense core, according to a scheme that we describe in Chapter 13. (Image from Bech and Atalay, [50].)
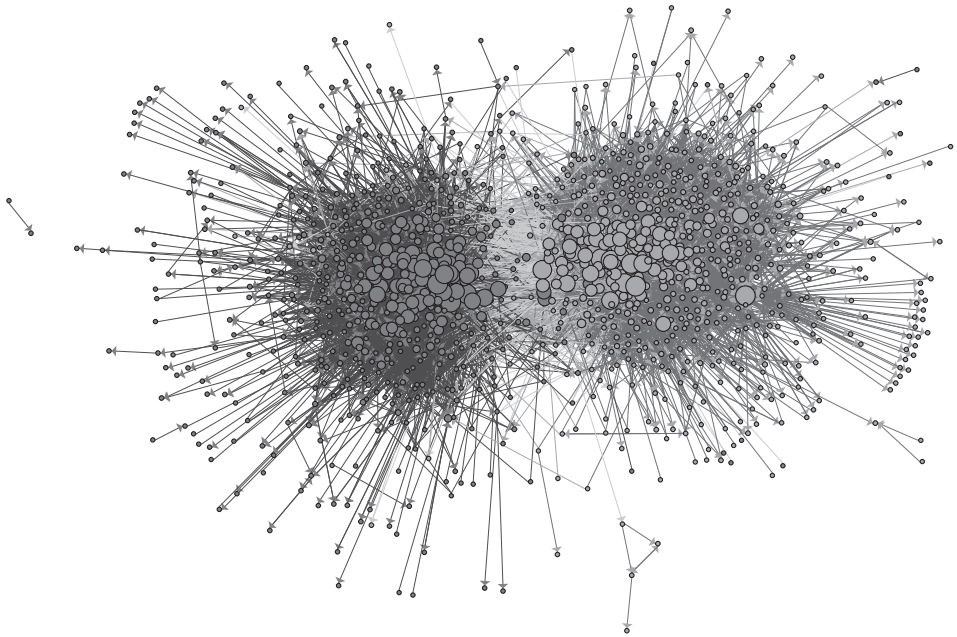
Figure 1.4. The links among Web pages can reveal densely knit communities and prominent sites. In this case, the network structure of political blogs prior to the 2004 U.S. presidential election reveals two natural and well-separated clusters [5]. (Image from Association for Computing Machinery, Inc.; http://www-personal.umich.edu/ ladamic/img/politicalblogs.jpg.)

Simply from their visual appearance, we can already see some of the complexity inherent in network structures. It is generally difficult to summarize succinctly the whole network; some parts are more or less densely interconnected, sometimes with central "cores" containing most of the links and sometimes with natural splits into multiple tightly-linked regions. Participants in the network can be more central or more peripheral; they can straddle the boundaries of different tightly-linked regions or sit squarely in the middle of one. Developing a language for talking about the typical structural features of networks is an important first step in understanding them.

Behavior and Dynamics. But the structure of the network is only a starting point. When people talk about the "connectedness" of a complex system, in general they are really talking about two related issues. One is connectedness at the level of structure – who is linked to whom – and the other is connectedness at the level of *behavior* – the fact that each individual's actions have implicit consequences for the outcomes of everyone in the system.

This means that, in addition to a language for discussing the structure of networks, we also need a framework for reasoning about behavior and interaction in network contexts. And just as the underlying structure of a network can be complex, so too can the coupled behavior of its inhabitants. If individuals have strong incentives to achieve good outcomes, then they not only will appreciate that their outcomes depend on how others behave, but they also take this into account in planning their own actions.

Figure 1.5. The rapidly growing popularity of YouTube is characteristic of the way in which new products, technologies, or innovations rise to prominence through feedback effects in the behavior of many individuals across a population. The plot depicts the number of Google queries for YouTube over time. The image comes from the site Google Trends (http://www.google.com/trends?q=youtube); by design, the units on the *y*-axis are suppressed in the output from this site.

As a result, models of networked behavior must take strategic behavior and strategic reasoning into account.

A fundamental point here is that, in a network setting, you should evaluate your actions not in isolation but with the expectation that the world will react to what you do. This means that cause-and-effect relationships can become quite subtle. Changes in a product, a Web site, or a government program can seem like good ideas when evaluated using the assumption that everything else will remain static, but in reality such changes can easily create incentives that shift behavior across the network in ways that were initially unintended.

Moreover, such effects are at work whether we are able to see the network or not. When a large group of people is tightly interconnected, these people often respond in complex ways that are only apparent at the population level, even though these effects may come from implicit networks that we do not directly observe. Consider, for example, the way in which new products, Web sites, or celebrities rise to prominence (as illustrated, for example, by Figures 1.5 and 1.6, which show the growth in popularity
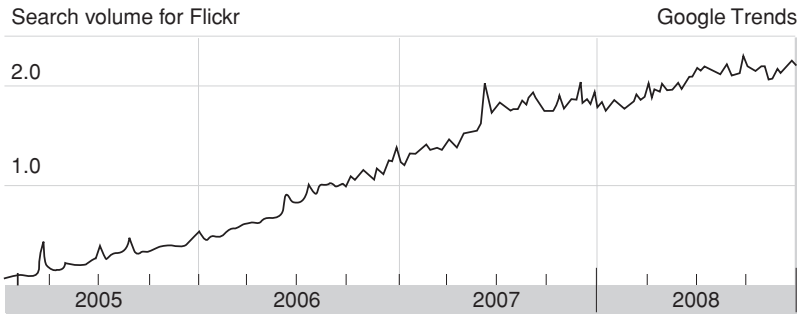

Figure 1.6. This companion to Figure 1.5 shows the rise of the social media site Flickr; the growth in popularity has a very similar pattern to that of other sites, including YouTube. (Image from Google Trends, http://www.google.com/trends?q=flickr.)

of the social media sites YouTube and Flickr, respectively, over the past several years). What we see in these figures is a growing awareness and adoption of a new innovation that is visible in aggregate, across a whole population. What are the underlying mechanisms that lead to such success? Standard refrains are often invoked in these situations: the rich get richer, winners take all, small advantages are magnified to a critical mass, and new ideas get attention that becomes "viral." But the rich don't always get richer and small advantages don't always lead to success. Some social networking sites flourish, like Facebook, while others, like SixDegrees.com, vanish. To understand how these processes work and how they are realized through the interconnected actions of many people, we need to study the dynamics of aggregate behavior.

A Confluence of Ideas. Understanding highly connected systems, then, requires a set of ideas for reasoning about network structure, strategic behavior, and the feedback effects they produce across large populations. These are ideas that have traditionally been dispersed across many different disciplines. However, in parallel with the increasing public interest in networks, there has been a coming-together of scientific fields around the topic of network research. Each of these fields brings important ideas to the discussion, and a full understanding seems to require a synthesis of perspectives from all of them.

One of the central goals in this book is to help bring about such a synthesis, by combining approaches that have traditionally been pursued separately. From computer science, applied mathematics, and operations research we draw on a language for talking about the complexity of network structure, information, and systems with interacting agents. From economics we draw on models for the strategic behavior of individuals who interact with each other and operate as members of larger aggregates. From sociology – particularly the more mathematical aspects concerned with social networks – we draw on a broad set of theoretical frameworks for talking about the structure and dynamics of social groups.

And the overall picture can help fill in pieces that are arguably missing from the intellectual landscape of each of these disciplines. Economics has developed rich theories for the strategic interaction among small numbers of parties, as well as for the cumulative behavior of large, homogeneous populations. The challenge is that much of economic life takes place in the complex spectrum between these extremes, with macroscopic effects that arise from an intricate pattern of localized interactions. Sociology has developed some of the fundamental insights into the structure of social networks, but its network methodology has been refined in the domains and scales where data collection has traditionally been possible – primarily, well-defined groups with tens to hundreds of people. The explosion of new contexts in which we find network data and network applications – including enormous, digitally mediated ones – leads to new opportunities for how we can pose questions, formulate theories, and evaluate predictions about social networks. Computer science, with the rise of the Web and social media, has had to deal with a world in which the design constraints on large computing systems are not just technological but also human – imposed by the complex feedback effects that human audiences create when they collectively use the Web for communication, self-expression, and the creation of knowledge. A fully satisfactory theory of network structure and behavior has the potential to address the simultaneous challenges encountered by all these fields.

A recurring theme underlying these challenges is the way in which networks span many different levels of scale and resolution. There are interesting questions that reach from the scale of small groups, such as the thirty-four–person social network in Figure 1.1, all the way up to the level of whole societies or economies, or to the body of global knowledge represented by the Web. In this book we examine networks both at the level of explicit structures, like those in Figures 1.1–1.4, and at the level of aggregate effects, like the popularity curves in Figures 1.5 and 1.6. As we look at networks of increasing scales, it becomes correspondingly more appropriate to take aggregate models into account. But the ability to work with massive network data sets has also enriched the picture, making it possible to study networks with billions of interacting items at a level of resolution where each connection is recorded. When an Internet search engine identifies the most useful pages from an index of the entire Web, for example, it is doing precisely this in the context of a specific task. Ultimately, it is an ongoing and challenging scientific problem to bridge these vastly different levels of scale so that predictions and principles from one level can be reconciled with those of others.

1.2 Central Themes and Topics

With this set of ideas in mind, we now introduce some of the main topics considered in this book and the ways in which these topics reinforce the underlying principles of networks. We begin with the two main bodies of theory that we will be building on: graph theory and game theory. These are theories of structure and behavior, respectively. Graph theory is the study of network structure, while game theory provides models of individual behavior in settings where outcomes depend on the behavior of others.

Graph Theory. In our discussion of graph theory, we focus particularly on some of the fundamental ideas from social network analysis, framing a number of graph-theoretic concepts in these terms. The networks in Figures 1.1 and 1.2 hint at some of these ideas. In the corporate e-mail communication network from Figure 1.2, for example, the communication is balanced between staying within small organizational units and cutting across organizational boundaries. This is an example of a much more general principle in social networks – that *strong ties*, representing close and frequent social contacts, tend to be embedded in tightly-linked regions of the network, whereas *weak ties*, representing more casual and distinct social contacts, tend to cross between these regions. Such a dichotomy suggests a way of thinking about social networks in terms of their dense pockets of strong ties and the ways in which they interact with each other through weaker ties. In a professional setting, it suggests a strategy for navigating one's way through the social landscape of a large organization by finding the *structural holes* between parts of the network that interact very little with each other. At a global scale, it suggests some of the ways in which weak ties can act as "shortcuts" that link together distant parts of the world, resulting in the phenomenon colloquially known as *six degrees of separation*.

Social networks can also capture the sources of conflict within a group. For example, latent conflicts are at work in the karate club social network from Figure 1.1. The people labeled 1 and 34 (the darker circles) are particularly central in the network of
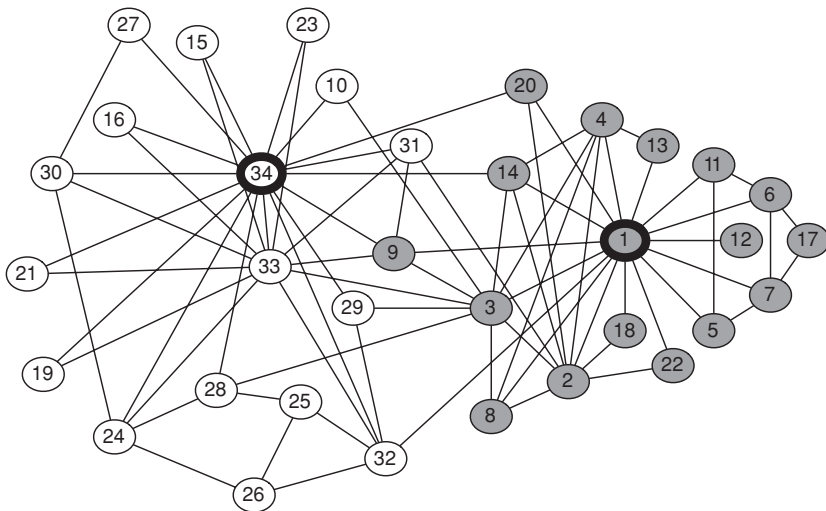
Figure 1.7. From the social network of friendships in the karate club from Figure 1.1, we can find clues to the latent schism that eventually split the group into two separate clubs (indicated by the two different shadings of individuals in the drawing).

friendships, with many connections to other people. On the other hand, they are not friends with each other, and in fact most people are only friends with one or the other of them. These two central people were, respectively, the instructor and the student founder of the club, and this pattern of noninteracting clusters was the most visible symptom of a conflict between them and their factions that ultimately splintered the group into two rival karate clubs, as shown in Figure 1.7. Later we will see how the theory of *structural balance* can be used to reason about how fissures in a network may arise from the dynamics of conflict and antagonism at a purely local level.

Game Theory. Our discussion of game theory starts from the observation that there are numerous settings in which a group of people must simultaneously choose how to act, knowing that the outcome will depend on the decisions made by all of them. One natural example is the problem of choosing a driving route through a network of highways at a time when traffic is heavy. For a driver in such a situation, the delays experienced depend on the pattern of traffic congestion arising not just from the driver's choice of route, but also from the choices made by all other drivers. In this example, the network plays the role of a shared resource, and the combined actions of its users can either congest this resource or use it more efficiently. In fact, the interactions among people's behavior can lead to counterintuitive effects; for example, adding resources to a transportation network can in fact create incentives that seriously undermine its efficiency, in a phenomenon known as *Braess's Paradox* [76].

Another example that will recur in several settings throughout the book is the problem of bidding in an auction. If a seller is trying to sell a single item using an auction, then the success of any one bidder in the auction (whether she gets the item, and how much she pays) depends not just on how she bids but also on how everyone else bids; an optimal bidding strategy should take this into account. Here too, counterintuitive effects

are at work: for example, if the seller introduces more aggressive pricing rules into the auction, he can make the strategic behavior of the bidders much more complex, and in particular induce optimal bidding that offsets whatever gains he might have expected to make from the new rules. Auctions represent a basic kind of economic interaction that we will generalize to more complex patterns of interactions in networks.

As a general part of our investigation of game theory, we abstract such situations with interdependent behavior into a common framework, wherein a collection of individuals must each commit to a *strategy*, thereby receiving a *payoff* that depends on the strategies chosen by everyone. Interpreting our preceding examples in this light, we see that the strategies available to a driver on a set of highways consist of the different options for routes he can take, and the payoff to this driver is based on his resulting travel time. In an auction, the strategies are the different choices for how to bid, and the payoff to a bidder is the difference between the value of the goods she receives and the price she pays. This general framework allows us to make predictions about how people will behave in a range of such situations. A fundamental part of this framework is the notion of *equilibrium* – a state that is "self-reinforcing" in that it provides no individual with an incentive to unilaterally change his or her strategy, even if that individual knows how others will behave.

Markets and Strategic Interaction in Networks. Once we have developed graph theory and game theory, we can combine them to produce richer models of behavior in networks. One natural setting for this exploration is in models of trade and other forms of economic activity. The interactions among buyers and sellers, or pairs of counterparties to a trade or loan, naturally forms a network. In Figure 1.3 we saw an example of such a network, with links between banks engaging in a loan. Figure 1.8 shows another example: a network representation of international trade among twenty-eight countries [262], in which the size of each country depicts its total amount of trade, and the thickness of each link connecting two countries indicates the amount of trade between them.

Where do these networks come from? In some cases, they are the traces of what happens when each participant seeks out the best trading partner it can find guided by how highly it values different trading opportunities. In other cases, they also reflect fundamental underlying constraints in the market that limit the access of certain participants to each other. In modern markets, these constraints could be institutional restrictions based on regulations; in other settings, they could be based on physical constraints like geography. For example, Figure 1.9 shows a map of trade routes in medieval Europe: when the physical movement of goods is costly and difficult, the economic outcome for different cities can depend significantly on where they are located in the underlying transportation network.

In all these settings, the network structure encodes a lot about the pattern of trade, and the success levels of different participants are affected by their positions in the network. Having a powerful position, however, depends not just on having many connections providing different options, but also on more subtle features – such as the power of the other individuals to which one is connected. Later we will see that this idea of network positions conferring power has been extended much more broadly and reaches beyond
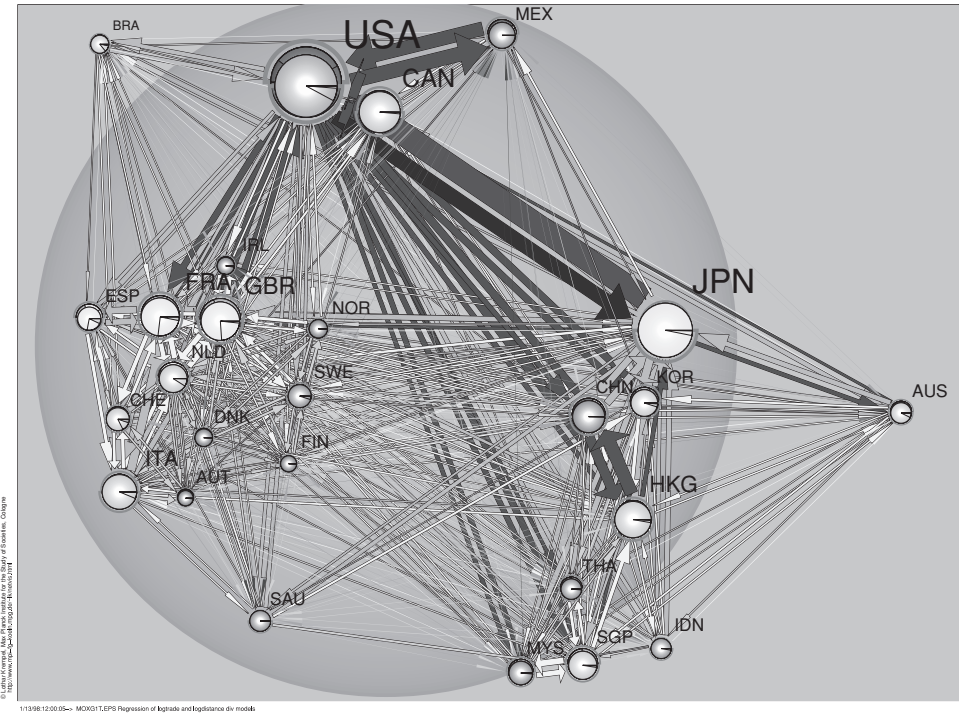
Figure 1.8. In a network representing international trade, one can look for countries that occupy powerful positions and derive economic benefits from these positions [262]. (Image from Carnegie Mellon University; http://www.cmu.edu/joss/content/articles/volume4/ KrempelPlumper.html.)

just economic exchange to suggest how power imbalances in many forms of social relationships may have their roots in the network patterns formed by the relationships.

Information Networks. The information we deal with online has a fundamental network structure. Links among Web pages, for example, can help us understand how these pages are related, how they are grouped into different communities, and which pages are the most prominent or important. Figure 1.4 illustrates some of these issues: it shows a network of links among political blogs constructed by Lada Adamic and Natalie Glance in the period leading up to the 2004 U.S. presidential election [5]. Although the network in the figure is too large to be able to see clearly the detailed structure around individual blogs, the image and its layout do convey the clear separation of the blogging network into two large clusters, which turn out to closely correspond to the sets of liberal and conservative blogs, respectively. From more detailed analysis of the raw linkage data underlying the image, it is possible to pick out the prominent blogs within each of these clusters.

Current Web search engines such as Google make extensive use of network structure in evaluating the quality and relevance of Web pages. To produce search results, these sites evaluate the prominence of a Web page based not just on the number of links it receives but on more subtle aspects of its position in the network. For example, a page