



Calhoun: The NPS Institutional Archive
DSpace Repository

NPS Scholarship

Publications

2008

Catching the "Network Science" Bug: Insight and Opportunity for the Operations Researcher

Alderson, D.L.

Alderson, D.L., 2008, Catching the "Network Science" Bug: Insight and Opportunity for the Operations Researcher. *Operations Research* 56(5): 1047-1065.
<https://hdl.handle.net/10945/36930>

defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

OR FORUM

Catching the “Network Science” Bug: Insight and Opportunity for the Operations Researcher

David L. Alderson

Operations Research Department, Naval Postgraduate School, Monterey, California 93943, dlalders@nps.edu

Recent efforts to develop a universal view of complex networks have created both excitement and confusion about the way in which knowledge of network structure can be used to understand, control, or design system behavior. This paper offers perspective on the emerging field of “network science” in three ways. First, it briefly summarizes the origins, methodological approaches, and most celebrated contributions within this increasingly popular field. Second, it contrasts the predominant perspective in the network science literature (that abstracts away domain-specific function and instead focuses on graph-theoretic measures of system structure and dynamics) with that of engineers and practitioners of decision science (who emphasize the importance of network performance, constraints, and trade-offs). Third, it proposes *optimization-based reverse engineering* to address some important open questions within network science from an operations research perspective. We advocate for increased, yet cautious, participation in this field by operations researchers.

Subject classifications: networks/graphs: theory; philosophy of modeling; engineering.

Area of review: OR Forum.

History: Received February 2007; revisions received October 2007, November 2007, December 2007; accepted January 2008.

Participate in the OR Forum discussion at <http://orforum.blog.informs.org>.

1. Introduction

Recent attention on the large-scale structure of many vital network systems has led to the proliferation of new theories that attempt to explain, predict, and control network behavior and evolution. The ubiquity of the network paradigm across many important and practical applications—including the Internet and communication systems, manufacturing systems and supply chains, national infrastructures, military systems, global markets, and social organizations—has created significant interest in whether there exist universal properties of networks that may be discovered and then applied to understand and manage them. To empower operations researchers looking to capitalize on these research trends, this paper provides a review and commentary about the potential benefits and pitfalls of recent approaches to “complex networks.”

As documented in a 2006 National Research Council (NRC) report (2006), a new research field called “network science” is focused on an interdisciplinary view of complex network systems. The NRC Report describes progress in this field and summarizes efforts to establish network science as an academic discipline. The scientific literature over the past several years (as measured by the quantity of publications) has emphasized phenomenological descriptions of these systems based on graph-theoretic properties and the interpretation of large-scale system measurements as the likely outcomes of random processes. For example, the application of statistical mechanics to graph theory

emphasizes the prevalence of universal statistical features, such as power laws, in the measurement, modeling, and assessment of network structure and behavior (e.g., Albert and Barabási 2002).

Broadly, the scientific questions of interest to researchers in network science include the following:

- Does there exist a *network structure* that is responsible for large-scale properties in complex systems? Typically, the properties of interest range from traditional engineering concepts such as performance and reliability, to opaque notions such as flexibility, adaptability, and sustainability.
- Are there *universal laws* governing the structure and behavior of complex networks? In particular, to what extent is *self-organization* (i.e., coordination from the “bottom up”) responsible for the emergence of system features not explained from a reductionist (i.e., “top down”) viewpoint?
- How can one assess the *vulnerabilities* or *fragilities* inherent in complex networks to avoid “rare, yet catastrophic” disasters (e.g., the August 14, 2003 power outage in the northeastern United States)? More practically, how should one design, organize, build, and manage complex networks?

Although in its infancy, network science has captured the interest of scientists, managers, policymakers, and the military. This is due in large part to the wide availability of academic and tutorial material at all levels. For example, there are survey papers (Barabási et al. 1999, Albert

and Barabási 2002, Newman 2003, Watts 2004), technical handbooks for students and practitioners (Baldi et al. 2003, Bornholdt and Schuster 2003, Dorogovtsev and Mendes 2003, Pastor-Satorras and Vespignani 2004, Ben-Naim et al. 2004, Newman et al. 2006), and even popular science books (Barabási 2002, Watts 2003, Buchanan 2003, Ball 2004). Empowered by advances in information technology that support the large-scale collection, storage, and sharing of real network data, researchers have developed new analytic and empirical techniques to study complex networks. Accordingly, the number of research projects and publications in the field is growing dramatically.

There are considerable differences between the mainstream network science literature and operations research (OR) in assumptions, modeling, and methods of analysis. As discussed below, there is a tendency in the network science literature to abstract away domain-specific functions, and focus instead on graph-theoretic measures of structure and dynamics. In contrast, engineers and practitioners of decision science are typically driven by application data and emphasize performance, constraints, and trade-offs in the design or operation of networks. Not surprisingly, these differences have important implications for the application of each approach to network decision problems.

However, the NRC Report and general public discourse on network science lack the "OR perspective," despite the deep contributions of OR to the study of networks. OR has been largely ignored in the network science literature—an exception is the introductory chapter of the retrospective anthology by Newman et al. (2006) that cites Ahuja et al. (1993) and Nagurney (1993) as exemplars—with the result that scientists or analysts, who look to this expanding body of research to learn the latest tools and techniques for analyzing real systems, obtain a limited (and sometimes misguided) view of what "matters" for network structure and behavior.

The objectives of this paper are twofold: (1) to provide an entry point for the OR community to engage network science by briefly reviewing the origins, contributions, and trends in this field, and (2) to present a conceptual framework for contrasting network science with traditional OR and engineering. Hopefully, this broader perspective facilitates critical thinking in the "complex networks debate" and highlights opportunities for contribution from operations researchers.

This paper is organized as follows. Section 2 presents a framework for the study of complex systems, comments on the challenges associated with complex network research, and highlights contributions in the study of networks within OR. Section 3 then reviews the origins, recent trends, and most celebrated results in network science and summarizes its academic impact. Section 4 presents a contrasting view of network science that incorporates notions of

design and optimization and highlights some major differences between network science and traditional engineering approaches. Specifically, we use the router-level Internet as a case study to illustrate the use of *optimization-based reverse engineering* as an alternative approach to the systematic investigation of network structure and function. Section 5 discusses the role of design in complex network systems, and §6 concludes by highlighting opportunities for contribution. Ultimately, this paper cautiously advocates for greater involvement in network science on the part of operations researchers, and it identifies a path for increased participation.

2. Networks as Complex Systems

A central challenge in the study of complex systems is understanding the relationship between system structure and function. For simplicity, we define *system structure* to mean the system *components* and their *interactions*, as well as the *constraints* and *uncertainties* governing them. *System function* then means the purposeful behavior resulting from that structure. For many "everyday" complex systems (e.g., economies, social organizations, living organisms), function must be inferred by approaching the system as an *artifact*. When such a system can be represented as a network, the network scientist will use observation, theory, and experiment to characterize its behavior and to infer the purpose for its structural features. The need to solve this *inverse problem*, that is, answering how the observed structure supports the perceived function, differs from the perspective of an engineer who presumes a well-defined notion of function and then approaches system structure with the intent of controlling the system or designing it from scratch.

For scientists across disciplines, the network paradigm has become popular for representing the interactions among discrete system components or as a discrete approximation to many continuous phenomena. The appeal of network models is that the mathematical tools and techniques apply, at least in principle, to any system representable as a graph. An important distinction in this paper is the difference between a *graph* (i.e., the mathematical object composed of vertices and edges) and a *network*, which consists of a graph *plus some data* (Ahuja et al. 1993, p. 33). This distinction is important because many complex-systems researchers view the domain-specific details as incidental to the development of elegant and abstract graph models, whereas the operations researcher typically seeks to employ the application-specific data that supplements a graph. In practice, however, the term "network" often lacks precise meaning and (like the term "system") serves as little more than a Rorschach test—allowing individuals to see the structural and behavioral patterns that are most familiar to them. The term "complex network" is even more ambiguous, despite its frequent use in many disciplines,

and we will not attempt a formal definition except to say that it is usually a network system with (1) a large number of components (complexity of size), (2) intricate relationships among components (complexity of interconnection), or (3) many degrees of freedom in the possible actions of components (complexity of interaction). Consequently, it is increasingly difficult (particularly to researchers who may have a limited view of network models and applications) to understand when different network-modeling approaches are appropriate.

Determining which aspects of the problem are essential and which can be safely abstracted away is a key question in developing an appropriate model of any system. The study of complex networks is no different, but is complicated sometimes by the stark differences in assumptions and methods that researchers from diverse backgrounds employ. It is critical to recognize that, despite the desire to obtain a universal view of complex networks, the results obtained from any particular domain are heavily influenced by its underlying perspective, and in extreme cases it is possible that the approaches taken by different researchers lead them to *opposite conclusions about one and the same system*. For example, Albert et al. (2000) use models of graph connectivity to claim that the Internet is vulnerable to attacks on the most highly connected routers, but Doyle et al. (2005) later show that a more realistic view of Internet structure and function reveals the network to be quite robust to attacks on highly connected routers, but vulnerable to hijacking of software protocols (something abstracted away from models based solely on graph connectivity). Thus, one must exercise caution when applying results from network science to decision problems, with particular scrutiny directed at the assumptions underlying the problem formulation and solution.

The use of graphs and networks as a framework for modeling combinatorial, operational, and structural problems predates recent interest in network science. The study of graphs in mathematics is attributed to Euler (1736) and the so-called Königsberg bridge problem, an instance of what is now known as the *postman problem* (see, for example, Evans and Minieka 1992, Chapter 8). Driven by applications in transportation, economics, electrical theory, and molecular theory, the study of graphs progressed until the early twentieth century, at which point one can identify the first network studies in what might be considered operations research. The economists Tolstói (1930), Kantorovich (1939), Hitchcock (1941), and Koopmans (1947) studied the implications of network structure for optimal resource allocation in production and transportation problems (see Schrijver 2002 for a discussion of this early history).

The study of networks by operations researchers grew with the development of linear programming (Dantzig 1948) and its application to problems in transportation (Dantzig 1951) and scheduling (Dantzig and Fulkerson

1954). From here, the use of networks in operations research proceeded in several directions. Considerable effort was directed at optimization aspects of networks, with Dantzig (1962) focused on simplex-based methods and Ford and Fulkerson (1962) focused on primal-dual combinatorial algorithms. Ahuja et al. (1993) document this and more recent history with over 150 applications of network flow problems. A key theme in this body of work is the special structure that a network provides for the development of extremely fast optimization algorithms.

Another related field of OR emphasizes user-driven models of economic equilibrium in complex network systems. Nagurney (2003) reviews this line of research that dates back to Quesnay (1758) and Cournot (1838). A key distinction here is the difference between *user optimization* and *system optimization*, and again, transportation problems were of particular importance (e.g., Beckmann et al. 1956). This theory of network dynamics and equilibria is now well documented (e.g., Florian and Hearn 1995, Giannessi and Maugeri 1995, Daniele 2006), and has been applied to a variety of systems including transportation networks (e.g., Ran and Boyce 1996), financial networks (e.g., Nagurney 2003), and supply chains (e.g., Nagurney 2006). A key idea here is that the structure and behavior of many complex network systems results from interacting decision processes between disparate agents, and understanding the way in which they "solve" coordinated problems via cooperation and/or competition is an active area of research (e.g., Johari et al. 2005, Acemoglu and Ozdaglar 2007). This type of problem is particularly difficult in a network context, where the agents often interface in a decentralized and asynchronous manner, and where the interaction of "selfish" agents often leads to suboptimal outcomes for the system as a whole (e.g., the so-called "price of anarchy" as summarized in Roughgarden 2005).

A vast operations research literature now exists on the application of network theory to a variety of decision problems. Table 1 summarizes recent activity within the INFORMS community, both by publication and application area. INFORMS journals do not represent a complete list of OR publications, and the categories used in this table are not exact, but Table 1 clearly illustrates that networks pervade this literature. Moreover, the prevalence of network-related problems addressed by recent Edelman Award winners and finalists (see <http://www.scienceofbetter.org/Edelman> for details) demonstrates the impact of OR in solving real-world, complex network problems.

Despite this long tradition in the use of network models by operations researchers and the wide availability of technical handbooks on network models in OR (e.g., Ball et al. 1995), it is network science that is having a considerable impact on scientists who are drawn to the study of complex networks. At the same time, the general popularity of network science is also showing signs of influencing decision makers at all levels. This may be reason enough for

Table 1. Two views into recent network research activity within the INFORMS community.

	2001	2002	2003	2004	2005	2006	2007*	Total
Recent activity, by publication								
<i>Management Science</i>	4	7	8	12	18	15	14	78
<i>Operations Research</i>	4	14	12	8	14	9	12	73
<i>Transportation Science</i>	13	9	8	12	10	10	4	66
<i>INFORMS Journal on Computing</i>	7	4	3	9	6	7	7	43
<i>Interfaces</i>	8	3	1	7	7	6	2	34
<i>Organization Science</i>	4	5	3	4	6	6	6	34
<i>Mathematics of Operations Research</i>	2	3	1	7	2	1	7	23
<i>Information Systems Research</i>	4	3	0	3	1	2	3	16
<i>Marketing Science</i>	0	1	1	2	3	1	2	10
<i>Manufacturing & Service Operations Management</i>	0	1	1	1	2	2	1	8
<i>Decision Analysis</i>	0	0	0	0	2	0	0	2
Total	46	50	38	65	71	59	46	387
Recent activity, by application area								
Mathematics: theory, computation	13	17	8	12	12	8	18	88
Business, management	6	10	9	13	13	6	12	69
Transportation, transit systems	15	6	6	13	13	10	2	65
Organizations, social systems	2	5	5	3	9	6	13	43
Manufacturing, production planning, supply chains	1	3	3	8	9	13	3	40
Data networks, telecommunications	5	5	2	2	5	9	6	34
Scheduling, delivery, assignment	1	1	1	6	5	4	1	19
Queueing, stochastic networks	2	2	3	6	3	1	2	19
Critical infrastructure protection	0	0	1	0	1	2	1	5
Military applications	0	1	0	0	1	0	0	2
Biomedical applications	0	0	0	2	0	0	0	2
Finance	1	0	0	0	0	0	0	1
Total	46	50	38	65	71	59	46	387

Notes. *These statistics are as recorded by *INFORMS Online* on October 1, 2007. A search of the term "network" in the title or abstract returned a total 387 entries.

operations researchers to pay attention to the trends in this new field of research.

3. The "New" Science of Networks

What is network science? The NRC report (National Research Council 2006, p. 3) concedes that "different research communities give different answers to [this] question," but goes on to assert that "network science is distinct from both network technology and network research: It is characterized by the discovery mode of science rather than the invention mode of technology and engineering." The report later adds, "network science consists of the study of network representations of physical, biological, and social phenomena, leading to predictive models of these phenomena."

Such a broad definition leads one to this question: What exactly is novel here? We defer the answer to the network science literature itself.

The title of this section comes from the introduction to a recent anthology of key network science papers as compiled by Mark Newman, Albert-László Barabási, and Duncan Watts—arguably three of the most recognized authorities in this field. The unmistakable double meaning

in their use of "new" is that the recent efforts to understand complex networks have departed from traditional approaches. Specifically, they claim (Newman et al. 2006, p. 4) that network science

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.

Although perhaps accurate when viewed from the lens of graph theory, this perspective does not reflect the application-driven research in OR that has been ongoing for more than 50 years.

An important issue in network science relates to the dynamic nature of networks, specifically, the distinction between *dynamics on networks* (i.e., behavior on top of a fixed graph structure) and *dynamics of networks* (i.e., the evolution of the graph structure itself) as noted by Watts (1999a). Of course, many phenomena of practical interest involve the interaction of the two. For example, in a

metabolic network, the activation of a gene may alter the biochemical pathways that in turn can alter other genes, and so on. In contrast, the tripping of a circuit breaker in an electrical grid may shift the current to other portions of the network, which in turn may trip other circuit breakers, further shifting the load and possibly leading to a *cascading failure*. Finally, the progression of a virus within a population may depend both on the properties of the disease it causes as well as the dynamics of the social network through which it is transmitted. Such complex behaviors are of primary interest in network science, and understanding these dependencies as well as their impact on system behavior is a key objective of the field. Newman et al. (2006, p. 4) further advocate the network science view as follows:

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... 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.

Despite this stated focus on network dynamics beyond applied graph theory, much of the recent work in network science seeks to characterize the connectivity of complex network systems.

3.1. Random Graphs as a Foundation

The structure of many important complex network systems is not known with certainty, either because it is not possible to inspect the networks directly or because the networks' large size and scope preclude a vantage point from which complete information can be obtained. For example, because administrative control of the Internet was given over to commercial entities in 1995, network owners and operators have stopped sharing topology information for proprietary and privacy reasons. Subsequent growth in the Internet's technologies and organizational entities has yielded a landscape where it is nontrivial even to visualize the network (Cheswick et al. 2000). In such cases, a primary challenge is to characterize system structure. Recent advances in information technology make it easier to measure, collect, and share empirical data about networks, but the fundamental issue is how to interpret and model relevant network features. For the Internet and many other complex systems, one popular approach has been to start with models based on *random graphs*.

The formal study of random graphs was popularized through the pioneering work by Erdős and Rényi (1959). Perhaps their most widely known model is one in which, for a given set of *vertices* (equivalently, *nodes*), one adds an *edge* (equivalently, *arc* or *link*) between each vertex pair with uniform probability p ($0 \leq p \leq 1$). Thus, for small values of p the graph is likely to be very sparse, and for large values of p the graph is likely to be dense, with the entire graph forming a single connected cluster. One of the more celebrated features of this model is that the overall connectivity of the graph undergoes a *phase transition* at a critical value γ , where for values of $p < \gamma$ the graph is likely to be broken into many small connected components, and for values $p > \gamma$ most of the nodes in the graph will almost surely belong to a single *giant component* (for a comprehensive review, see Bollobás 1998). That this phenomena is reminiscent of phase transitions in physics has made random graphs a popular starting point for researchers familiar with statistical mechanics.

Random graphs have been a popular starting point for modeling large network systems for which only connectivity properties matter (or are available for study). In the context of the Internet, the first popular network topology generator to be used for the simulation of Internet protocols was the model by Waxman (1988), which is a variation of the classical Erdős-Rényi random graph in which nodes are connected according to a nonuniform probability that is inversely proportional to the distance between them. The rationale for this model is the observation that long-distance links are expensive and thus unlikely to be used in practice. The Waxman model was later abandoned in favor of other models that explicitly generate nonrandom structure (see Li et al. 2004 for a review of this history), but the point is that, in the absence of domain-specific details, random graphs have served as a natural "null hypothesis" for evaluating properties of network structure.

A popular approach to testing this null hypothesis has been to compare the measured connectivity features of real networks with those of random graphs. Two features have received the most attention: *power-law statistics* and *small-world phenomena*.

Power-Law Statistics. When the distribution of *degree* (i.e., number of connections, denoted here as x) for each node is appropriately represented in the tail by a function $d(x) \propto cx^{-\alpha}$, where $\alpha > 0$ and c is a positive finite constant, then one says that the network exhibits a *power-law* (or equivalently, a *scaling distribution*). In contrast, the degree distribution for random Erdős-Rényi type graphs follows the form of a Poisson variable, specifically, $d(x) = e^{-(N-1)p}((N-1)p)^x/x!$ in the limit as the number of nodes $N \rightarrow \infty$ (Newman et al. 2002), thus making these types of graphs unrealistic representations for graphs exhibiting this power-law phenomenon.

Power laws have been observed for more than a century within the social sciences and economics (income

distributions, city populations), linguistics (word frequencies), ecology (the size and frequency of forest fires), biology (the distributions of species within plant genera and mutants in old bacterial populations), molecular biology (cellular metabolism and genetic regulatory networks), and the Internet (router graphs and the World Wide Web); see Mitzenmacher (2004) and Li et al. (2006) and references therein for details. Newman (2005) provides a comprehensive review of the mathematics and mechanisms underlying power laws. To the extent that these systems can be modeled using some type of network, these examples lend evidence to arguments in favor of power laws as universal features in many complex network structures.

Small-World Phenomena. Recent attempts to understand the structure of large social networks has shown that many naturally occurring or man-made systems have certain statistical features that make them look simultaneously *regular* (in the sense of a lattice) and *random* (in the sense of an Erdős-Renyí graph). As first documented by Watts and Strogatz (1998), these graphs are characterized succinctly by three statistics: *characteristic path length* is the average shortest number of edges between connected pairs of distinct vertices; *average vertex degree* is the average number of incident edges to a vertex; and *clustering coefficient* is the (dimensionless) frequency with which three connected vertices are fully connected (i.e., they form a triangle). For two graphs of equal size and having the same average vertex degree, random graphs tend to have lower characteristic path lengths when compared to regular graphs. Conversely, random graphs tend to have lower clustering values when compared to regular graphs. However, there is an intermediate class of graphs that has relatively high clustering coefficients and short characteristic path lengths. In the context of social networks, this signature characterizes the "small-world phenomenon"—the seemingly frequent experience by which two strangers learn that they share a common acquaintance or are similarly "connected" through a short sequence of individuals. Empirical studies report that small-world features also exist outside social networks: in the Internet, road networks, electric power grids, food chains, and neural networks (Watts 1999b). This ubiquity has generated interest in small worlds as universal models of complex networks.

3.2. A "Physics View" of Networks

Much of network science has employed tools, techniques, and a mindset from physics—the usual approach abstracts away the domain-specific details of a problem to isolate and investigate its most "essential" features. When applied to large-scale networks, the standard view has been to combine the use of graph theory with the tools and techniques of statistical mechanics (Barabási et al. 1999, Albert and Barabási 2002, Newman 2003, Amaral and Ottino 2004). In particular, one typically treats the network as a member of a *random ensemble* and then often models its evolution as a

dynamical system, governed by (differential) equations and with an emphasis on equilibrium behavior. This approach has enabled the development of some elegant mathematical tools, such as mean-field models for networks (Newman et al. 2000), with the caveat that each result implicitly relies on key assumptions underlying the chosen method for analysis (e.g., the network is sufficiently large scale and homogeneous).

The use of random ensembles to model network structure ties in naturally with random graph theory, and it has opened the world of networks to a large community of researchers trained in statistical mechanics. The result has been an explosion in descriptive models that attempt to characterize the structure and evolutionary dynamics of graphs, often with random graphs as the underlying null hypothesis for comparison. Power laws have received particular emphasis in this context because the traditional statistical physics perspective views power-law distributions as evidence of an internal self-sustaining critical state, often associated with a phase transition (Bak 1996, Ball 2004). In the face of phenomena that cannot be explained by "traditional" models (e.g., Erdős-Renyí graphs), this approach focuses on specialized models that reproduce and thereby "explain" the observed emergent behavior (Bak 1996, Barabási 2002, Buchanan 2003, Ball 2004).

Scale-Free Networks. A recently popular model used to explain the apparent ubiquity of power laws in network structure is the so-called *scale-free network (SFN)*. Originally introduced by Barabási and Albert (1999), the use of "scale-free" comes from their observation that "many large random networks share the common feature that the distribution of their local connectivity is free of scale, following a power law" (p. 510). This definition has never been made precise (see the commentary in Bollobás and Riordan 2003), and the resulting ambiguity has created confusion about the applicability of scale-free network models (for details, see Li et al. 2006). In essence, scale-free network models argue that the power laws observed in many complex networks are the large-scale result of simple random processes that occur during network evolution. Thus, scale-free networks follow naturally from other models inspired by statistical physics, including *self-organized criticality (SOC)*; see Bak (1996) and *edge-of-chaos (EOC)*; see Kauffman (1993). In all cases, the generation mechanisms in these models are generic and independent of system-specific details. They assume that interactions are essentially random, but have some macroscopic statistic tuned to a special point, such as a bifurcation point (EOC), a critical density (SOC), or a power-law degree distribution (SFN).

The simplest method for generating a scale-free network is via *preferential attachment*, in which (1) the network grows by the sequential addition of new nodes, and (2) each newly added node is more likely to connect with a node that already has many connections. Formally, a newly added node connects to an existing node k

with probability $\Pi(k) \propto (d_k)^\beta$, where d_k is the degree of node k (in contrast to traditional random graph models, where $\Pi(k) = p$ for all k , i.e., $\beta = 0$). As a consequence, high-degree nodes are likely to get more and more connections (a phenomenon also known as “the rich get richer” or the “Matthew effect”), and the end result is a power law in the distribution of node degree. By tuning β , one can achieve a wide range of power laws consistent with those observed in real networks (Albert and Barabási 2002). One can also generate random graphs with specified degree distributions (e.g., Aiello et al. 2000). Because many empirically observed power laws are consistent with the statistics produced by these *degree-based* network models, scale-free network structure is argued to be universal (Barabási 2002).

The proposed structure of scale-free networks resulting from degree-based generation has serious implications for any system it represents. Perhaps most critical is the advertised presence of highly connected central *hubs* (representing the highest-degree nodes) that yield a “robust yet fragile” connectivity structure. That is, the scale-free topology is simultaneously robust to the random loss of nodes (giving the network “error tolerance”), but fragile to targeted worst-case attacks (causing “attack vulnerability”). This latter feature, when applied to the Internet, has been termed its “Achilles’ heel” (Albert et al. 2000), implying that targeted attacks on the highest-connectivity nodes could destroy its overall connectivity and cripple its performance. Bollobás and Riordan (2003, 2004) provide treatment of scale-free graphs from a random graph perspective.

Researchers have also used scale-free networks to model sexual contact networks (Liljeros et al. 2001), and the application of scale-free models to both Internet and social networks advertises important implications for the understanding of virus propagation—either computer viruses in the Internet or infectious diseases in social networks—because the presence of highly connected central hubs makes scale-free networks highly susceptible to epidemic outbreaks (Pastor-Satorras and Vespignani 2001). This research suggests that the solution to epidemics is to target vaccination and prevention strategies at these central hubs, whether they be highly connected Internet nodes (Briesemeister et al. 2003) or highly connected individuals within a social network (Dezsö and Barabási 2002, Pastor-Satorras and Vespignani 2002).

Small-World Networks. In parallel to the characterization of the small-world phenomenon, Watts and Strogatz (1998) demonstrate that this statistical signature can be reproduced by relatively simple graph models that interpolate between regular and random graph structures. The simplest model is one in which a d -dimensional square lattice consisting of nearest-neighbor connections is rewired or supplemented with a relatively few, random “shortcut links”—reducing the overall average path length without changing the relatively high clustering. Chung and Lu

(2003) provide complimentary treatment of small-world graphs from a classical random graph perspective.

The study of this and other features for small-world networks has been largely conducted using statistical physics. For example, Newman and Watts (1999a) show that the number of shortcut links needed to obtain the small-world effect behaves according to a phase transition. Their model is a d -dimensional lattice of size N in each dimension (thus having a total N^d vertices) with nearest-neighbor edge connections and periodic boundary conditions (i.e., for $d = 1$, the lattice is a ring). With this model, they show that when additional shortcut connections are added in a uniformly random manner according to probability p , the model undergoes a phase transition or crossover (moving from a “small-world regime” to a “large-world regime”) as p approaches zero. They calculate the exact value of the single critical exponent for the system (Newman and Watts 1999a, b) and also develop a solution for the average path length and for the distribution of path lengths (Newman et al. 2000). In addition, Newman and Watts (1999b) consider *percolation* (a popular framework in statistical mechanics; see Stauffer and Aharony 1992 for background) on these small-world graphs as a simple model of disease transmission in a social network. Using a setup in which each vertex is “infected” with probability ρ , they identify when ρ leads to the formation of a giant component of infected vertices (intended to represent the epidemic threshold). Calloway et al. (2000) later extend this to include the possibility of either link or node “failures” in networks having general degree distributions.

The small-world model has been used to represent many types of social networks, including collaboration networks (Newman 2001), trust networks (Gray et al. 2003), and community structure (Girvan and Newman 2002). However, the ability of this framework to capture a seemingly universal statistical signature has led to an even more prolific use of this model outside of social networks. Small-world models have been used as models of general communication networks (Comellas et al. 2002), as models of file-sharing communities (Jovanović et al. 2001, Iamnitchi et al. 2004), and models of the Internet (Jin and Bestavros 2002). In the context of biological systems, small-world models have been used to represent neural networks (Bohland and Minai 2001), chemical reaction networks (Gleiss et al. 2001), and metabolic networks (Wagner and Fell 2001).

The observation that many of the same networks, such as collaboration networks and the Internet, can be classified as both scale-free and having the small-world property has led to model extensions that blur their distinction (e.g., Klemm and Eguiluz 2002 propose variations on preferential attachment mechanisms in scale-free models that increase clustering similar to the small-world phenomena). Amaral et al. (2000) argue that scale-free networks are a subclass of small worlds, along with *broad-scale* networks (having

Table 2. Growth in the "network science literature" by publication area.

	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007*	Total
"High impact"	1	1	5	4	17	13	22	16	9	4	92
Physics	1	7	26	62	124	139	230	260	350	286	1,485
Biology, chemistry, medicine	0	1	4	16	22	31	67	80	94	77	392
Computer science	0	1	2	7	10	22	47	61	64	19	233
Sociology and economics	0	1	2	6	7	11	14	22	15	16	94
Complex systems	0	1	1	2	3	7	11	13	18	22	78
Engineering	0	0	1	2	7	4	13	15	22	12	76
Applied mathematics	0	0	0	0	2	6	6	10	29	21	74
Earth science	0	1	1	2	7	4	6	11	11	0	43
Business and management	0	0	0	1	2	1	4	6	9	1	24
Total	2	13	42	102	201	238	420	494	621	458	2,591

Notes. *These statistics are as recorded by the *Web of Science* October 1, 2007. A search of the terms "scale-free" or "small-world" returned 3,151 entries, from which 560 were irrelevant to network science. Here, "High impact" includes the journals *Nature*, *Science*, *Proceedings of the National Academy of Sciences of the U.S.A.*, *Scientific American*, and *American Scientist*. Because the *Web of Science* only lists publications in peer-reviewed journals, scientific communities where a majority of the publications appear in conferences (e.g., computer science) or as working papers (e.g., complex systems) are most likely underrepresented in this table.

a truncated power-law distribution) and *single-scale* networks (having an exponential type of degree distribution). Although this work has provided a taxonomy of graph structures, it has also contributed to an environment where both scale-free graphs and small-world graphs are applied universally to any complex network bearing the appropriate statistical signature.

3.3. Scientific Impact

Network science is much broader than the study of scale-free and small-world systems, yet we emphasize these topics here because they are two of the most prominent and celebrated subjects. Also, their development provides historical context for the ongoing work that is now appearing regularly across a diversity of scientific communities. Despite its short history, network science is having considerable impact on the way that complex network systems are viewed and studied. Although it is difficult to measure directly the impact of a scientific movement, it is possible to quantify scientific activity in terms of the number of publications and citations on particular topics, such as scale-free and small-world networks. Table 2 shows the yearly publication activity by discipline in this network science literature. The most vigorous activity has been in the physics journals, with biology and computer science also growing in recent years.

The literature on scale-free and small-world networks is only a subset of the ongoing work on complex network systems. Nonetheless, these two models have been extremely influential, as indicated by Table 3, which lists the most highly cited articles. Remarkably, the top 10 publications have received well over 10,000 citations, suggesting that the impact of network science is large. Although articles on scale-free and small-world networks have not been prominent in the INFORMS journals, there is growing interest in complex network systems within the community (e.g.,

Management Science presented a special issue on complex systems across disciplines in July 2007).

3.4. Criticism of Network Science

The application of network science to practical problems has been met with considerable skepticism. A basic criticism of network science is that by reducing a complex network to a simple graph, one eliminates all of the key features that differentiate one system from another. Some of the strongest criticism has come in the context of biology, where a proper accounting of biological details in the context of small-world graphs (Arita 2004) and scale-free graphs (Tanaka 2005) shows previous applications to have yielded specious results. Keller (2005) provides a particularly sharp critique of scale-free graphs as they pertain to biological systems. Another popular area of application for network science has been the Internet, and here again it has been shown that ignoring the presence of heterogeneous components, layered architectures, and feedback dynamics can lead to serious misinterpretation of observed graph structure (Doyle et al. 2005). Specifically, Li et al. (2006) demonstrate that evidence for the "Achilles' heel" vulnerability of the router level of the Internet is an artifact of the inappropriate application of random ensemble models and has no relevance to the actual network. Although there is evidence suggesting that the Internet is indeed "robust, yet fragile," this fact has nothing to do with any perceived scale-free structure (Doyle et al. 2005).

A second argument against current approaches in network science is that the almost exclusive emphasis on statistical characterizations of graph structure causes the following practical problems.

1. Many statistical descriptions do not uniquely characterize the system of interest, and there often exists considerable diversity among graphs that share any particular statistical feature. This is particularly true for scale-free networks, e.g., recent work by the author and his colleagues

Table 3. Top 25 most highly cited publications in the “network science literature.”

Rank	Article	Times cited
1	Watts, D. J., S. H. Strogatz. 1998. Collective dynamics of “small-world” networks. <i>Nature</i> 393 (668).	2,244
2	Barabasi, A. L., R. Albert. 1999. Emergence of scaling in random networks. <i>Science</i> 286 (543).	2,110
3	Albert, R., A. L. Barabasi. 2002. Statistical mechanics of complex networks. <i>Rev. Modern Phys.</i> 74 (1).	1,972
4	Newman, M. E. J. 2003. The structure and function of complex networks. <i>SIAM Rev.</i> 45 (2).	960
5	Jeong, H., B. Tombor, R. Albert, Z. N. Oltval, A. L. Barabasi. 2000. The large-scale organization of metabolic networks. <i>Nature</i> 407 (6804).	903
6	Strogatz, S. H. 2001. Exploring complex networks. <i>Nature</i> 410 (6825).	884
7	Albert, R., H. Jeong, A. L. Barabasi. 2000. Error and attack tolerance of complex networks. <i>Nature</i> 406 (6794).	747
8	Dorogovtsev, S. N., J. F. F. Mendes. 2002. Evolution of networks. <i>Adv. Phys.</i> 51 (4).	636
9	Giot, L., J. S. Bader, C. Brouwer, A. Chaudhuri, B. Kuang, et al. 2003. A protein interaction map of <i>Drosophila melanogaster</i> . <i>Science</i> 302 (5651).	550
10	Milo, R., S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, U. Alon. 2002. Network motifs: Simple building blocks of complex networks. <i>Science</i> 298 (5594).	489
11	Amaral, L. A. N., A. Scala, M. Barthelemy, H. E. Stanley. 2000. Classes of small-world networks. <i>Proc. Nat. Acad. Sci. USA</i> 97 (21).	475
12	Ravasz, E., A. L. Somera, D. A. Mongru, Z. N. Oltvai, A. L. Barabasi. 2002. Hierarchical organization of modularity in metabolic networks. <i>Science</i> 297 (5586).	457
13	Pastor-Satorras, R., A. Vespignani. 2001. Epidemic spreading in scale-free networks. <i>Physical Rev. Lett.</i> 86 (14).	440
14	Tong, A. H. Y., G. Lesage, G. D. Bader, H. M. Ding, H. Xu, et al. 2004. Global mapping of the yeast genetic interaction network. <i>Science</i> 303 (5659).	412
15	Barabasi, A. L., R. Albert, H. Jeong. 1999. Mean-field theory for scale-free random networks. <i>Physica A</i> 272 .	364
16	Newman, M. E. J. 2001. The structure of scientific collaboration networks. <i>Proc. Nat. Acad. Sci. USA</i> 98 (2).	352
17	Cohen, R., K. Erez, D. ben-Avraham, S. Havlin. 2000. Resilience of the Internet to random breakdowns. <i>Physical Rev. Lett.</i> 85 (21).	308
18	Liljeros, F., C. R. Edling, L. A. N. Amaral, H. E. Stanley, Y. Aberg. 2001. The web of human sexual contacts. <i>Nature</i> 411 (6840).	280
19	Newman, M. E. J., S. H. Strogatz, D. J. Watts. 2001. Random graphs with arbitrary degree distributions and their applications. <i>Physical Rev. E</i> 64 02(2).	275
20	Girvan, M., M. E. J. Newman. 2002. Community structure in social and biological networks. <i>Proc. Nat. Acad. Sci. USA</i> 99 (12).	261
21	Newman, M. E. J., D. J. Watts. 1999. Scaling and percolation in the small-world network model. <i>Physical Rev. E</i> 60 (6).	221
22	Pastor-Satorras, R., A. Vazquez, A. Vespignani. 2001. Dynamical and correlation properties of the Internet. <i>Physical Rev. Lett.</i> 87 (25).	217
23	Wagner, A., D. A. Fell. 2001. The small world inside large metabolic networks. <i>Proc. Roy. Soc. London Ser. B</i> 268 (1478).	189
24	Barahona, M., L. M. Pecora. 2002. Synchronization in small-world systems. <i>Physical Rev. Lett.</i> 89 (5), Art. No. 054101.	188
25	Newman, M. E. J. 2001. Scientific collaboration networks: I. Network construction and fundamental results. <i>Physical Rev. E</i> 64 01(1)	183

Note. These statistics are as recorded by the *Web of Science* on October 1, 2007.

(Li et al. 2006) has shown there is enough diversity among graphs having the *same* power-law node degree distribution that, although indistinguishable by this parsimonious characterization, these graphs can actually be interpreted as “opposites” when measured against other performance-based metrics. Figure 1 shows a simple example of four graphs that have the *same* degree sequence, which happens to be heavy tailed. A problem with many popular approaches to generating graphs using random ensembles is that these methods are more likely to yield graphs that look like Figure 1(d), with highly structured graphs like those in Figure 1(a)–(c) appearing so rarely as to be effectively ignored altogether (see Alderson and Li 2007). Many of the celebrated results for scale-free graphs stem from a belief that the presence of a power law in the node degree distribution of a graph *necessarily* implies a network structure qualitatively similar to Figure 1(d), a belief that is incorrect.

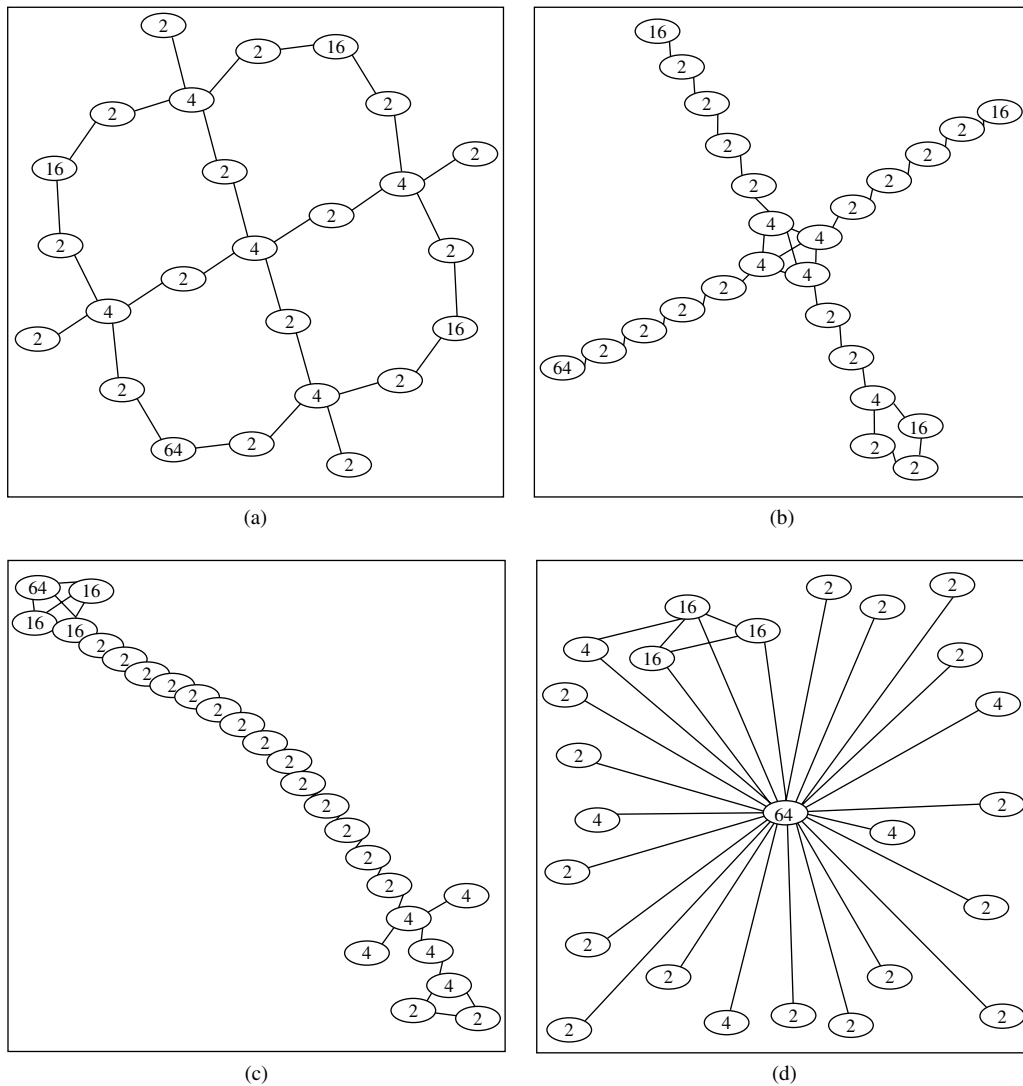
2. Because many processes can generate similar graphs, one can infer little about the underlying processes that caused an observed feature. More generally, network science has been accused of producing merely descriptive, not explanatory, models (Willinger et al. 2002).

3. The blind application of small-world and scale-free models wherever their statistical signatures are found creates a danger for researchers not familiar with the underlying or implicit assumptions of these models. Watts himself warns that

claiming that everything is a small-world network or a scale-free network not only oversimplifies the truth but does so in a way that can mislead one to think that the same set of characteristics is relevant to every problem (Watts 2003, p. 304).

At the core of the criticism toward network science is its *applicability* to real problems. Mitzenmacher (2006) casts

Figure 1. Four graphs with the same degree sequence but with obvious structural differences.



Notes. The label on each node indicates its total degree. Degree-one nodes have been omitted for visual clarity.

this criticism in the context of the following natural progression of published scientific results: (1) Observe, (2) Interpret, (3) Model, (4) Validate, and (5) Control. He states (p. 527),

most research on power laws [and perhaps network science in general] has focused on observing, interpreting, and modeling, with a current emphasis on modeling. As a community, we have done almost nothing on validation and control, and we must actively move towards this kind of research.

In other words, it is now time to shift the emphasis in network science research toward the development and validation of explanatory models of network structure and function, and it is in this area that the OR community has an important role to play.

4. A Contrasting Approach to Complex Networks

Whereas the previous discussion highlighted the most celebrated topics in network science, this section offers a more subjective view of the importance of engineering and OR in the study of complex networks. The intent is to contrast the existing network science approach with a perspective that instead emphasizes system performance, resource constraints, and design trade-offs as essential.

4.1. An Engineering View of System Structure and Function

The engineering approach to complex systems follows a different paradigm from network science. In engineering, any notion of system function must be well defined (perhaps specified a priori), and *forward engineering* is the

process by which one explores the relationship between system structure and function to design the components and interactions that ensure desired behavior. However, for many real systems the notion of function is not really understood, is often subject to interpretation, and is rarely defined in any formal sense. This ambiguity makes the direct application of forward engineering (e.g., via optimization) to the study of network science somewhat awkward because a well-posed mathematical formulation is typically not available from the outset.

Network science fits more naturally with *reverse engineering*, defined here as the process by which one models system structure to explain the observed function. Reverse engineering is critical in the design of many complex technologies, but it is less prominent in traditional OR. One point of contact with reverse engineering in the optimization literature is the concept of *inverse optimization* by Ahuja and Orlin (2001). Adopting their framework for a linear programming problem of the form $\min(cx: x \in X)$ and a feasible point $x_0 \in X$, the inverse problem is $\min(\|d - c\|: d \in \text{Inv}(x_0, X))$, where $\text{Inv}(x_0, X) = \{d: x_0 \text{ optimizes the math program } \min(dx: x \in X)\}$. In other words, one seeks the cost vector d that is “closest” to the original vector c using an appropriate definition of distance (e.g., $\|d - c\| = \sum_i |c_i - d_i|$) such that the feasible point x_0 is an optimal solution to this modified mathematical program. Ahuja and Orlin (2001) demonstrate several useful relationships between an optimization problem and its inverse (e.g., if the original problem is an LP, then so is its inverse) and develop solutions for the inverse minimum cost spanning tree, the inverse minimum cost flow, and the inverse minimum cut problems (Ahuja and Orlin 2002).

The primary question for operations researchers in this context is whether or not the structure and function of a complex network can be interpreted as the result of some (possibly implicit) optimization process. The power of an optimization-based approach to a complex system structure has been documented in several contexts. Carlson and Doyle (1999) introduce the notion of *highly optimized tolerance (HOT)* to demonstrate how highly variable event sizes (i.e., power laws) in systems optimized by engineering design can arise as the result of trade-offs between yield, resource costs, and risk tolerance. They argue that the ubiquity of power-law phenomena in the natural and man-made world may simply be the result of an inherent drive for systems to improve their performance while adhering to constraints imposed by scarce resources, physical limitations, or a hostile environment. They assert that *robustness* (i.e., the maintenance of some desired system characteristics despite uncertainties in the system’s components and/or environment) in complex systems is a constrained and limited quantity that must be diligently managed. In their view, most complex systems of interest are *highly optimized* in the sense that performance and behavior objectives are achieved by highly structured, rare, nongeneric system configurations that arise from iterative design either in natural

systems (via evolution) or man-made systems (via engineering). In turn, the characteristics of these HOT systems are high performance, highly structured internal complexity, yet apparently simple and robust external behavior, with the potential for rare but catastrophic cascading failures initiated by small perturbations (Carlson and Doyle 2002).

Fabrikant et al. (2002) present the first explicit attempt at using the HOT concept for network modeling and generation under the title of *heuristically optimized trade-offs*. They propose a model of network access design based on incremental growth that optimizes a trade-off between the local connection cost and the overall distance to other nodes in the network. More specifically, they consider a process in which each new node i is connected to the existing network according to the solution of $\min_{j < i} \alpha \cdot \text{dist}(i, j) + h_j$, where $\text{dist}(i, j)$ is the distance between nodes i and j , and where h_j measures the “centrality” (e.g., the average number of hops to other nodes in the network) of node j . They show that changing the relative weight α of these two terms in the overall objective function yields a spectrum of topological structures, with the resulting node degree distributions ranging from exponential (nonheavy tailed) to scaling (heavy tailed). Berger et al. (2003) later showed the claim of strict scaling for the heavy-tailed case to be incorrect (i.e., the resulting degree distribution follows a power law only up to a cutoff), but this is not relevant here. Although this work illustrated the power of optimization-based formulations to yield heavy-tailed distributions in topology generation, its construction was not intended as a model of real networks.

Can the objectives of network science be addressed using optimization-based reverse engineering? The problem in practice is that the types of networks under consideration are rarely as clean as the linear programs in Ahuja and Orlin (2001), and one still faces the challenge of having to choose from among an almost endless list of system properties the few features that are most relevant. However, recent research demonstrates that the application of inverse optimization provides insight into the structure and function of some complex networks, including the Internet.

4.2. Case Study: The Router-Level Internet

The Internet may be the most important complex network of the past decade, and its increasing presence and importance in daily life make it a popular object of study in the network science literature. The Internet has been shown to exhibit both scale-free and small-world properties (Adamic 1999, Barabási and Albert 1999, Pastor-Satorras and Vespignani 2002), and has inspired many of the high-profile discoveries in network science (e.g., the “Achilles’ heel” of scale-free networks by Albert et al. 2000). Although the multilayered architecture of the Internet protocol stack means that there is no single representation for the Internet as a network (Alderson et al. 2006), one network of practical importance is the *router-level* Internet, in which nodes represent routers and links between

nodes represent one-hop connectivity by the Internet Protocol (IP). The structure of the router-level Internet has practical implications for network provisioning, protocol performance, and system reliability (for example, tolerance of router loss resulting from failure or attack).

Most efforts in network science to model the router-level Internet have focused on matching observed connectivity statistics, typically power laws (Li et al. 2004 review these *degree-based* models). As noted, this approach suffers because of the inherent diversity among graphs having the same degree distribution. An alternate approach (Alderson et al. 2003) to router-level topology is to consider the technological and economic factors affecting the decisions of Internet Service Providers (ISPs) in the construction and provisioning of router-level networks. The argument is that only by considering the domain-specific details of real networks can one move beyond descriptive representations to develop explanatory models that reflect the causal forces driving their evolution.

A natural means to incorporate domain-specific details is to use reverse engineering. Consider a general mathematical program for traditional engineering design: given a definition of system performance $f(x)$ and a feasible region $X = \{x: g(x) \geq 0, h(x) = 0\}$ (both possibly nonlinear), find the best system design given by $x^* = \arg \max\{f(x): x \in X\}$. The reverse-engineering problem is then as follows: given a working system (i.e., feasible point x^0), find the objectives and constraints such that the structure produces the function (e.g., find f, X such that x^0 is a "good" solution to $\max\{f(x): x \in X\}$). Although this type of inverse optimization problem does not follow Ahuja and Orlin (2001) exactly, it shares the same basic form.

This type of inverse problem is underconstrained, making domain knowledge essential to narrow the possible choices for valid solutions. For the router-level Internet, first principles suggest *system throughput* as a reasonable design objective and *router technology* as an important constraint on the feasible region for possible designs. Specifically, because a router can only process a finite number of packets per time unit, there is an inherent trade-off between the number of connections a router can support (i.e., its degree) and the amount of traffic that can be sent on those connections (i.e., the *bandwidth* of each connection). In the simple case where all routers are equal, a router with more connections can only support lower bandwidths. This type of *bandwidth-degree* constraint defines a simple, but effective, feasible region for router-level design, and this perspective provides the means to interpret the results from various empirical studies as feasible points (i.e., the x^0).

Using router throughput constraints and a realistic model of traffic demand, Li et al. (2004) generate networks via heuristic optimization that provide high throughput by placing the highest-degree nodes toward the network periphery for traffic aggregation purposes. When evaluated with the same constraints and traffic demands, degree-based

models of equal size and having the same degree distribution have poor throughput characteristics because their highest-degree hubs (which typically reside in the center of the network) serve as bottlenecks. Different choices in the objective function and/or constraints yield different measures of performance and feasibility, but here the emphasis is on finding a parsimonious representation of the drivers of network evolution and not on a system that is formally optimal. Additional validation against empirical data for real networks (Alderson et al. 2005) shows that these optimization-based models not only capture structural features of router-level graphs not found in their degree-based counterparts, but they also complement ongoing empirically based efforts to reverse-engineer the Internet.

Whereas network science emphasizes graph connectivity and generating random ensembles to identify the "most likely" model that fits observation, the approach here leverages different assumptions and yields sharply different results. As reported by Doyle et al. (2005), high-degree routers in the Internet must be toward the network periphery (where they enable traffic aggregation) and not in the network core (where attacking them could fragment the network, as reported by Albert et al. 2000). The use of inverse optimization in this context stems from an assumption that the observed system has a specialized structure that has "evolved" (e.g., via iterative design) to achieve some system objective. This starting assumption gives the approach both its strengths and weaknesses.

4.3. Pros and Cons of a Reverse-Engineering Approach

Considerable effort remains to develop systematic reverse-engineering techniques for complex systems, but the example above shows how an optimization-based framework may capture key tensions and trade-offs in the evolution of some networks. Moreover, reverse engineering via optimization offers several advantages over approaches based primarily on graph-theoretic characterizations.

Pro 1: *Reverse engineering takes direct advantage of domain-specific details that differentiate the network under study from its generic underlying graph.* Focusing on the domain-specific objectives and constraints for a particular network system ensures a minimal level of realism. For example, emphasis on network throughput and technology constraints reveals that router-level networks generated from degree-based methods typically either cannot be built from existing equipment or have such poor relative performance that they would never be implemented in practice.

Pro 2: *By capturing the tensions and trade-offs in the construction of complex networks, the reverse-engineering approach potentially provides insight into the decisions faced by network owners, operators, and designers.* For example, an optimization-based approach to ISP network design and operation provides a natural context for investigating the relationship between decisions about network

provisioning, traffic engineering, and demand estimation (Alderson et al. 2006).

Pro 3: *A successful reverse-engineering effort invites a straightforward impact assessment from potential changes to a problem’s objectives and constraints.* For example, how could new technologies affect the design decisions of ISPs in building and operating their networks? This form of sensitivity analysis is typically not possible with existing network science approaches.

Pro 4: *An optimization-based formulation can accommodate additional empirical observations, system constraints, or objectives.* A potential problem with models intended to reproduce aggregate statistics (e.g., node degree distributions) is that the discovery of new graph-theoretic signatures often requires considerable model redesign. Matching aggregate statistics is only secondary evidence of successful optimization-based reverse engineering, so this approach is robust to changes in modeling emphasis when new or competing graph descriptions are discovered.

Pro 5: Finally, *reverse engineering often provides the opportunity to study important related problems.* For example, optimization-based reverse engineering of the Internet’s topology and protocols has led researchers to consider the extent to which the entire Internet protocol stack can be interpreted as a giant resource allocation problem, with individual protocols solving particular optimization sub-problems in a decentralized, asynchronous manner (Chiang et al. 2007).

In essence, by focusing on optimization as a *modeling process*, not a specific modeling outcome (i.e., the solution to any one optimization problem), one can systematically study how particular objectives and constraints shape the large-scale structure and behavior of complex networks. With this perspective, optimization-based reverse-engineering approaches such as HOT serve best as a conceptual framework (or a modeling methodology), not a specific model for complex networks.

Despite its potential advantages over existing techniques, optimization-based reverse engineering must overcome several challenges to be appropriate and successful.

Con 1: Optimization-based reverse engineering by itself will not identify a parsimonious representation of essential system features. However, it does provide a means to systematically test how different objectives and/or constraints translate to different outcomes in network behavior.

Con 2: Large network problems are hard, and in practice they are often solved only heuristically. In such cases, any “solution” is a result of not just the problem formulation (objective and constraints) and the problem data (parameter values), but also *the approximation technique itself* (Alderson et al. 2003). This significantly complicates the use of inverse optimization.

Con 3: Reverse engineering is unlikely to reproduce an existing complex system in exact detail. For example, focusing on the technological and economic forces shaping the decisions of the ISP may not reproduce the existing

Internet, but the hope is to find “realistic, yet fictitious” models to use when real networks are not available for proprietary or security reasons (Alderson et al. 2005).

Con 4: It is possible that real decisions affecting the design of complex technological or social systems are neither consistent nor rational, and thus do not fit this mathematical formulation. For example, anecdotal evidence from ISP operators suggests that *what ought to be done* is often very different from *what was done* in the construction and operation of real systems.

In other words, the inexact nature of any underlying optimization problem means that in practice it may be difficult either to isolate the primary objectives and constraints in some complex systems or to validate them against measurements from real systems. More fundamentally, it remains uncertain what role, if any, design plays in the formation of many complex networks.

5. The Role of Design in Complex Networks

The use of optimization as a means to explore the relationship between complex network structure and function assumes that *design* in some form—possibly implicit, decentralized, heuristic, or ad hoc—plays a role in the evolution of the system. In contrast, much of the complex systems and network science literatures emphasize *emergent phenomena* and is focused on understanding the simple, random processes that give rise to complex behaviors. This tension leads to a fundamental debate: *Are complex network systems the result of design?* Here, we briefly review three key underlying issues, while also highlighting the different views from engineering and network science.

Can Complex Networks Be Engineered? It has recently been argued that engineering is about the design and operation of systems that are complicated, but not complex (Ottino 2004). The distinction suggests that engineering systems are well understood and well behaved—Ottino’s example is a watch having thousands of parts, but whose behavior as a group can be understood a priori from established theory that allows one to compute the interaction and ultimate performance as a system (see also Amaral and Uzzi 2007). In contrast, Ottino (p. 399) notes that

The hallmarks of complex systems are adaptation, self-organization and emergence—no one designed the web or the metabolic processes within a cell.

From a traditional engineering perspective that emphasizes design in support of a well-defined function, this statement may be accurate. However, the design objectives of modern engineering are shifting from traditional notions such as performance, function, and efficiency to opaque notions such as flexibility, evolvability, and survivability. As a result, approaches such as reverse engineering are

becoming important for understanding the relationship between structure and function. With this enhanced perspective, the view of engineering expressed above may be too narrow.

The Internet's original architects did not conceive the current World Wide Web (WWW), but they did intend a network that would support diverse applications and that could change with the overall system needs (Clark 1988). Thus, an explicit goal was to design a network that could outlast the ability of its designers to specify what any particular application might do, and in this regard it is hard to argue with the genius of the current architecture. Also, it is crucial to distinguish the creative content of the WWW (i.e., webpages and hyperlinks) from the actual technology enabling that content (i.e., the hypertext transfer protocol, or HTTP). Engineers did not design the creative content, but they did design the software protocols and hardware enabling it.

The ongoing demand for individuals, devices, and information to be connected is changing the types of problems that engineers must solve in practice. For example, the business imperatives of many technology companies drive them to design, mass produce, and deploy Internet-enabled devices or software without a precise understanding of how they will behave when connected "in the wild." National, state, and local governments must invest in protection of critical infrastructure without complete knowledge of how the system components will respond in the presence of an accident, failure, or attack. The need to address uncertainty, not only in terms of model inputs and the operating environment but *in the system objectives themselves*, is already forcing engineering into the world of the complex. In the future, either engineers will need to be comfortable working on complex systems that lack succinct functional requirements, or engineering as a discipline will need to establish a new vocabulary for describing function in complex systems.

Self-Organization vs. Design. The existing complex systems perspective and traditional engineering have contrasting approaches to self-organization and the role of randomness. As noted, the application of statistical physics to network problems presumes that large-scale system structure and behavior can be understood in terms of random ensembles and their statistical properties, and it emphasizes the "most likely" graph features arising in the equilibrium of some proposed dynamics. A focal point has been explaining the emergent features that arise out of this inherently probabilistic setting, and the ubiquity of similar phenomena across systems serves as evidence of universal self-organization.

In engineering, self-organization is typically a *design objective*—that is, the desire to minimize the need for human intervention, such as self-configuration during system startup, self-adaptation to environment changes, or self-healing from component failures (Alderson and Willinger

2005). The "organized complexity" that results from efforts to create simplicity through the use of (often hidden) underlying system complexity is very different from the complexity typically studied in mainstream network science (Alderson and Doyle 2007). Moreover, the primary use of randomness in engineering models is to account for uncertainty *that needs to be managed*, not as a driver of system dynamics. Mixing inherent uncertainty with hard, system-specific constraints drives engineers toward "hand crafted" designs that are extremely rare from a traditional random ensemble perspective. Thus, the answers that network science and engineering each find typically occupy distinctly different, and often disjoint, regions of the overall space of possible system configurations.

The Significance of Power Laws. The (re)discovery of power laws has generated considerable interest and controversy, and here again the prevailing network science view of the world contrasts sharply with engineering. At the heart of the debate is the frequent association made by researchers trained in statistical physics between power laws, the critical state of a phase transition, and self-organization. Barabási (2002, p. 77) captures this notion eloquently when he writes that nature's normal abhorrence of power laws is suspended

if the system is forced to undergo a phase transition. Then power laws emerge—nature's unmistakable sign that chaos is departing in favor of order.

This view of power laws as exotic and unexpected phenomena has created great interest in the physics literature, where considerable effort has gone to cataloging the existence of power laws across a diversity of systems.

Engineers often care more about the heavy-tailed nature of power laws than their precise mathematical form. Heavy tails arise naturally in insurance (e.g., risk modeling), computer science (e.g., load balancing), and optimization (e.g., restart methods in combinatorial search). They are also ubiquitous in disaster data, describing losses in both deaths and dollars (CRED 2006). They are important because their mean behavior is typically meaningless (e.g., insurance losses are dominated by the "rare, but catastrophic" events), and managing the behavior of systems that encounter them is an open area of study. Also, there exist longstanding arguments by Bookstein (1990) and Mandelbrot (1997), suggesting that the strong invariance properties of power laws make them the natural null hypothesis for highly variable phenomena. Because power laws arise naturally by many mechanisms (Newman 2005), they may be considered "more normal than Normal," i.e., they should be no more surprising than Gaussian data (Willinger et al. 2004). This latter perspective rejects the need to develop special models that explain the ubiquity of power laws.

The issue of whether or not naturally occurring complex systems are the result of design in a traditional engineering sense may be a red herring. A better question may be whether or not one can find important *design elements* in

the structure of naturally occurring and man-made complex systems. If so, the key question becomes, *Are these design elements the result of an evolutionary process that systematically rewards "good" configurations while punishing "poor" ones?* If the answer to this question is affirmative, then this presence of feedback in the evolution of the system will make optimization-based reverse engineering an important research tool.

Biological systems are the most obvious instances of highly evolved systems, with the role of "master designer" played by natural selection. It is sometimes suggested that this is not actually design because specific configurations are found by the random processes of mutation and sexual recombination, because historical precedent (e.g., the "frozen accident") plays a key role, and because the resulting solutions are not truly optimal. However, the evidence of design in the artifacts themselves is unmistakable. A key difference between natural selection and engineering is that nature has had millions of years to search the design space of possible configurations, with billions of trials in each case. Also, the ultimate objective (i.e., survivability) is not currently well understood from an engineering design perspective. Although engineers also use an iterative process in building complex systems, they have far fewer resources at their disposal when searching for good designs. It is precisely the need to find "good" configurations under severe resource constraints that separates engineering from other disciplines.

What often appears to the outside observer as emergent self-organization can often be understood in terms of rigorous mathematics and engineering that explain the inherent "design" in many complex systems (Alderson and Willinger 2005). This is the case for structural features of the Internet's router-level topology. Recent work on the Transmission Control Protocol (TCP) and Active Queue Management (AQM) has also shown that these Internet protocols, largely the result of tinkering and intuition, can now be understood as primal-dual optimization algorithms solving a global resource allocation problem (e.g., Kelly et al. 1998, Kelly 2001, Low and Srikant 2004). In contrast to previous arguments in favor of TCP behavior as a complex and chaotic phenomenon (Veres and Boda 2000, Solé and Valverde 2001), the reverse engineering of a rigorous mathematical framework has demonstrated why the existing protocols have worked well in the past, and it now suggests how to design their next-generation improvements (e.g., Wei et al. 2006). Thus, new approaches to network engineering are rising to address the challenges posed by complex networks, but considerable work remains.

6. A Path Forward

The strength and weakness of network science depends on the answer to the question: *What meaningful conclusions can one draw about a system based solely on its underlying network structure?* In some cases, the answer may be "a

great deal" and in others "not very much." Just as combinatorics enrich graph theory, network dynamics yields interesting models for graph formation and evolution. However, for decision makers interested in the study of real complex networks, it is clear that there should be much more to network science.

6.1. A Need to Study "Organized Complexity"

Sixty years ago, Warren Weaver (then director of natural sciences of the Rockefeller Foundation in New York City) coined the term "disorganized complexity" to refer to the types of systems particularly suited for the application of statistical mechanics (Weaver 1948). His example is that of billiard balls, for which classical dynamics provide exact descriptions of a small number of balls interacting on a table, but where the computational requirements for tracking a large number of balls becomes burdensome. In this context, the power of statistical mechanics is that, for a giant table consisting of millions or billions of interacting balls, one can answer with precision certain questions related to average properties of the system. However, Weaver (pp. 537–538) pointedly warns that:

the methods of statistical mechanics are valid only when the balls are distributed, in their positions and motions, in a helter-skelter, that is to say a disorganized, way. For example, the statistical methods would not apply if someone were to arrange the balls in a row parallel to one side rail of the table, and then start them all moving in precisely parallel paths perpendicular to the row in which they stand. Then the balls would never collide with each other nor with two of the rails, and one would not have a situation of disorganized complexity.

In other words, nonrandom organization in the structure of a system—a scenario that Weaver termed "organized complexity"—can render the tools of statistical mechanics inappropriate.

Should the scenario described by Weaver, that of carefully arranged billiard balls, be a concern to researchers studying complex systems? Many researchers answer "no." They argue that within the ensemble of all possible arrangements of billiard balls, the configurations described by Weaver are so rare as to constitute a set of measure zero. However, such an argument implicitly relies on an assumption that all such configurations are feasible and perhaps even equally likely. For a simple system of billiard balls, this assumption may be appropriate. Recently, however, the mathematical models originally conceived to describe systems such as billiard balls are later adopted as representations for other systems, such as systems of interacting people, computers, vehicles, cells, or genes. In these systems, all configurations are *not* feasible, simply because survival for these systems means performing a particular function or achieving a particular task, and not all configurations do so.

Arguments in favor of organized complexity assert that Weaver's example of carefully arranged billiard balls is

prevalent throughout "highly evolved" systems and thus should be central to the study of complex systems. More specifically, evolution and/or engineering design (which build upon "good" configurations that achieve required function) in these systems essentially use feedback to take the structure of complex systems very far away from "average" or "most likely" configurations. For high-technology and biological applications, the organization found in real-world systems carefully supports the required function, and their "designed" nature means that their structures are necessarily nonrandom and will be adequately represented by random processes only rarely.

How does one assess whether or not the complexity of a particular network system is organized or disorganized? One simple approach is to ask: *What is the effect of arbitrary perturbations to network structure that change certain aspects of its connectivity while leaving others invariant?* For example, because scale-free networks claim to take their properties primarily from their node degree distribution, then arbitrary rewiring that preserves this distribution is not believed to disrupt the network's most essential properties. However, for the router-level Internet, such changes can negatively impact network throughput by several orders of magnitude (Li et al. 2004). Similarly, in metabolic processes claimed to be scale-free, rewiring destroys all cellular function (Tanaka 2005).

We propose the following heuristic to test for organized or disorganized network complexity:

- When a network is sufficiently homogeneous such that its connectivity can be arbitrarily rewired to preserve its large-scale statistics without disrupting its functionality, and when domain-specific features outside the model can be ignored or treated as uniformly random, then we conjecture that this type of network is amenable to the tools and techniques of disorganized complexity.
- When a network has evolved through feedback (either by iterative design or via some form of natural selection), when the domain-specific features outside the model are important and/or highly evolved, or when arbitrary rewiring destroys its functionality (even when the overall statistics do not change), then we suggest the need to study the system as one having organized complexity.

Alderson and Doyle (2007) contrast these notions of complexity as applied to complex engineering systems, including critical infrastructures. The main idea is that the need for complex function by many naturally occurring and man-made systems results in an organization (or "design") for which the tools of disorganized complexity (e.g., statistical mechanics) are inappropriate.

Perhaps as poignant as his recognition for the need to study organized complexity is Weaver's (1948, pp. 540–541) insight into the necessary tools for doing so:

Science must, over the next 50 years, learn to deal with these problems of organized complexity. Is there any promise on the horizon that this new advance can really be accomplished? ... Out of [World War II] have come two new

developments that may well be of major importance in helping science to solve these complex twentieth-century problems. The first piece of evidence is the wartime development of new types of electronic computing devices. ... The second of the wartime advances is the "mixed-team" approach of operations analysis.

The study of organized complexity remains nascent, but the tools, techniques, and past contributions of OR make it well suited to address the questions posed by network science.

There is opportunity for researchers who can blend the aspirations of network science with the need to solve real, practical decision problems about the design, operation, and management of complex networks. Although some level of abstraction or simplification may always be required, ignoring the domain-specific features of real systems creates serious pitfalls. One seeks balance in model realism: to examine key tensions and trade-offs in the large-scale interaction of network components while still respecting the role of architecture, dynamics, and feedback in system behavior.

An important open topic relates to the *drivers of network formation*. Most of network science has focused thus far on the "what" as it pertains to relevant network structure and the "how" in terms of the possible causes of that structure, with little attention paid to "why" the network was formed in the first place. In other words, *what is the problem that is being solved by the network?* The answer to the "why" underlying network formation is of paramount importance for reverse-engineering efforts, yet the progress to date is relatively uneven across disciplines. For highly evolved technological or biological systems, the answer can be conceptually simple: reinforcement of what achieves desired function (and thus confers advantage). Reverse engineering uses this starting point to explore the relationship between system structure and function.

Research on social and economic networks has paid particular attention to the drivers of network formation, with emphasis on human incentives, a complexity not currently addressed in the mainstream network science literature. More specifically, economic theory suggests that networks form because the individual agents (nodes) derive some utility from the connectivity, either individually or as a social whole (e.g., Jackson 2006). Similarly, the formation of networks in social systems is often attributed to notions of social capital and embeddedness (see Borgatti and Foster 2003 for a review and typology of network models in organizational research). An enhanced understanding of the drivers of social network formation in the context of complex system structure and function would go a long way to answering several important questions. For example, can the organizational structure of the modern corporation be viewed as the solution to some type of design problem, and if so, what is it that the corporation is designed to achieve? Can the organizational structure of terrorist networks or military dictatorships be interpreted as a rational solution to a particular design problem involving extreme

constraints? If so, does relaxing these constraints offer a better alternative to undermining their functionality than direct attempts at network interdiction?

6.2. Conclusion

If publication trends are an accurate reflection of scientific activity, network science will continue to be a popular topic across disciplines. Although network science has been successful in capturing the attention and imagination of researchers, managers, and policymakers, considerable work remains before we will attain the required proficiency to predict, control, and design complex network systems. Most of the existing work has focused on descriptive approaches to network dynamics (i.e., the “what” and the “how” of network formation and growth) as well as the implications for dynamical behavior on top of these networks. In comparison, relatively little progress has been made in the development of explanatory models for network structure and function (i.e., the “why” underlying network formation), and this shortcoming often creates a sharp disconnect in the application of network science to real systems. This article contributes a first step toward bridging this gap by suggesting optimization-based reverse engineering as a systematic approach to the study of complex network structure and function.

For more than half a century, the OR community has been quietly solving some of the most challenging problems related to the practical design, operation, and management of networks exhibiting “organized complexity.” However, when it comes to the research agenda now popularized by network science, OR has been an underutilized resource, with the result that many decision makers tasked with important problems are headed in a direction that does not benefit from this vast body of theory and experience. Is network science simply a fad, something that will soon enough fade? It is too soon to tell, but in the meantime, increased participation, critical thinking, and leadership on the part of our community can only improve the level of understanding and quality of decisions being made in networks of all kinds.

Acknowledgments

The author thanks John Doyle and Walter Willinger for ongoing conversations regarding organized complexity; Jeffrey Rothal for library assistance; and Richard Larson for first suggesting this topic as an article for OR Forum. This manuscript benefited greatly from the feedback of numerous NPS colleagues, two anonymous referees, and a helpful area editor. This work was supported, in part, by ONR MURI Award #N0001408WR20242 and an NPS institutional grant.

References

Acemoglu, D., A. Ozdaglar. 2007. Competition and efficiency in congested markets. *Math. Oper. Res.* **32**(1) 1–31.

- Adamic, L. 1999. The small world web. *Proc. ECDL'99, Paris. Lecture Notes in Computer Science*, Vol. 1696. Springer, Berlin/Heidelberg, 443–452.
- Ahuja, R. K., J. B. Orlin. 2001. Inverse optimization. *Oper. Res.* **49**(5) 771–783.
- Ahuja, R. K., J. B. Orlin. 2002. Combinatorial algorithms for inverse network flow problems. *Networks* **40**(4) 181–187.
- Ahuja, R. K., T. L. Magnanti, J. B. Orlin. 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Upper Saddle Ridge, NJ.
- Aiello, W., F. Chung, L. Lu. 2000. A random graph model for massive graphs. *Proc. ACM STOC*, Portland, OR, 171–180.
- Albert, R., A.-L. Barabási. 2002. Statistical mechanics of complex networks. *Rev. Modern Phys.* **74**(1) 47–97.
- Albert, R., H. Jeong, A.-L. Barabási. 2000. Attack and error tolerance of complex networks. *Nature* **406** 378–382.
- Alderson, D. L., J. C. Doyle. 2007. Can complexity science support the engineering of critical network infrastructures? *Proc. IEEE Conf. Systems, Man, Cybernetics*, Montreal, Quebec, Canada, 44–51.
- Alderson, D., L. Li. 2007. Diversity of graphs with highly variable connectivity. *Phys. Rev. E* **75**. 046102.
- Alderson, D., W. Willinger. 2005. A contrasting look at self-organization in the Internet and next-generation communication networks. *IEEE Comm. Magazine* **43**(7) 94–100.
- Alderson, D., J. Doyle, R. Govindan, W. Willinger. 2003. Toward an optimization-driven framework for designing and generating realistic Internet topologies. *ACM Comp. Comm. Rev.* **33**(1) 41–46.
- Alderson, D., L. Li, W. Willinger, J. C. Doyle. 2005. Understanding Internet topology: Principles, models, and validation. *IEEE/ACM Trans. Networking* **13**(6) 1205–1218.
- Alderson, D., H. Chang, M. Roughan, S. Uhlig, W. Willinger. 2006. The many facets of Internet topology and traffic. *AIMS J. Networks and Heterogeneous Media* **1**(4) 569–600.
- Amaral, L. A. N., J. M. Ottino. 2004. Complex networks—Augmenting the framework for the study of complex systems. *Eur. Phys. J. B* **38** 147–162.
- Amaral, L. A. N., B. Uzzi. 2007. Complex systems—A new paradigm for the integrative study of management, physical, and technological systems. *Management Sci.* **53**(7) 1033–1035.
- Amaral, L. A. N., A. Scala, M. Barthelemy, H. E. Stanley. 2000. Classes of small-world networks. *Proc. Nat. Acad. Sci. USA* **97** 11149–11152.
- Arita, M. 2004. The metabolic world of *Escherichia coli* is not small. *Proc. Nat. Acad. Sci. USA* **101**(6) 1543–1547.
- Bak, P. 1996. *How Nature Works: The Science of Self-Organized Criticality*. Copernicus, New York.
- Baldi, P., P. Frasconi, P. Smyth. 2003. *Modeling the Internet and the Web: Probabilistic Methods and Algorithms*. John Wiley & Sons, Chichester, UK.
- Ball, P. 2004. *Critical Mass: How One Thing Leads to Another*. Farrar, Straus and Giroux, New York.
- Ball, M. O., T. L. Magnanti, C. L. Monma, G. L. Nemhauser, eds. 1995. *Network Models. Handbooks in Operations Research and Management Science*, Vol. 7. Elsevier, Amsterdam.
- Barabási, A.-L. 2002. *Linked: The New Science of Networks*. Perseus Publishing, Cambridge, MA.
- Barabási, A.-L., R. Albert. 1999. Emergence of scaling in random networks. *Science* **286** 509–512.
- Barabási, A.-L., R. Albert, H. Jeong. 1999. Mean-field theory for scale-free random networks. *Physica A* **272** 173–187.
- Beckmann, M. J., C. B. McGuire, C. B. Winsten. 1956. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT.
- Ben-Naim, E., H. Frauenfelder, Z. Toroczkai, eds. 2004. *Complex Networks. Springer Lecture Notes in Physics*, Vol. 650. Springer, Berlin.
- Berger, N., B. Bollobás, C. Borgs, J. T. Chayes, O. Riordan. 2003. Degree distribution of the FKP network model. *Proc. ICALP*, Eindhoven, The Netherlands, 725–738.

- Bohland, J. W., A. A. Minai. 2001. Efficient associative memory using small-world architecture. *Neurocomp* **38–40** 489–496.
- Bollobás, B. 1998. *Modern Graph Theory*. Springer-Verlag, New York.
- Bollobás, B., O. Riordan. 2003. Robustness and vulnerability of scale-free random graphs. *Internet Math.* **1** 1–35.
- Bollobás, B., O. Riordan. 2004. Coupling scale-free and classical random graphs. *Internet Math.* **1**(2) 215–225.
- Bookstein, A. 1990. Informetic distributions, part I: Unified overview & part II: Resilience to ambiguity. *J. Amer. Soc. Info. Sci.* **41**(5) 376–386.
- Borgatti, S. P., P. C. Foster. 2003. The network paradigm in organizational research: A review and typology. *J. Management* **29**(6) 991–1013.
- Bornholdt, S., H. G. Schuster, eds. 2003. *Handbook of Graphs and Networks: From the Genome to the Internet*. Wiley-VCH, Weinheim, Germany.
- Briesemeister, L., P. Lincoln, P. Porras. 2003. Epidemic profiles and defense of scalefree networks. *ACM Workshop on Rapid Malcode (WORM)*, Washington, D.C., 67–75.
- Buchanan, M. 2003. *Nexus: Small Worlds and the Groundbreaking Theory of Networks*. W. W. Norton, New York.
- Calloway, D. S., M. E. J. Newman, S. H. Strogatz, D. J. Watts. 2000. Networks robustness and fragility: Percolation on random graphs. *Phys. Rev. Lett.* **85**(25) 5468–5471.
- Carlson, J. M., J. C. Doyle. 1999. Highly optimized tolerance: A mechanism for power laws in designed systems. *Phys. Rev. E* **60** 1412–1428.
- Carlson, J. M., J. C. Doyle. 2002. Complexity and robustness. *Proc. Nat. Acad. Sci. USA* **99** 2539–2545.
- Cheswick, W., H. Burch, S. Branigan. 2000. Mapping and visualizing the Internet. *Proc. USENIX, Annual Tech. Conf.*, San Diego, 1–12.
- Chiang, M., S. H. Low, A. R. Calderbank, J. C. Doyle. 2007. Layering as optimization decomposition. *Proc. IEEE* **95** 255–312.
- Chung, F., L. Lu. 2003. The average distance in a random graph with given expected degrees. *Internet Math.* **1** 91–113.
- Clark, D. D. 1988. The design philosophy of the Darpa Internet protocols. *ACM Comp. Comm. Rev.* **18**(4) 106–114.
- Comellas, F., M. Mitjana, J. G. Peters. 2002. Broadcasting in small-world communication networks. C. Kalamanis, L. Kirousis, eds. *Proceedings in Informatics 13*. Carleton Scientific, Waterloo, Ontario, Canada, 73–85.
- Cournot, A. A. 1838. *Researches into the Mathematical Principles of the Theory of Wealth*. Macmillan, London. English translation, 1897.
- CRED: The Centre for Research on the Epidemiology of Disasters. 2006. EM-DAT: The OFDA/CRED international disaster database. Université Catholique de Louvain, Brussels, Belgium.
- Daniele, P. 2006. *Dynamic Networks and Evolutionary Variational Inequalities*. Edward Elgar Publishing, Inc., Cheltenham, UK.
- Dantzig, G. B. 1948. Programming in a linear structure. Technical report, Comptroller, United States Air Force, Washington, D.C.
- Dantzig, G. B. 1951. Application of the simplex method to the transportation problem. T. C. Koopmans, ed. *Activity Analysis of Production and Allocation*. John Wiley & Sons, New York, 359–373.
- Dantzig, G. B. 1962. *Linear Programming and Extensions*. Princeton University Press, Princeton, NJ.
- Dantzig, G. B., D. R. Fulkerson. 1954. Minimizing the number of tankers to meet a fixed schedule. *Naval Res. Logist. Quart.* **1** 217–222.
- Dezsö, Z., A.-L. Barabási. 2002. Halting viruses in scale-free networks. *Phys. Rev. E.* **65** 055103.
- Dorogovtsev, S. N., J. F. F. Mendes. 2003. *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press, Oxford, UK.
- Doyle, J. C., D. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, W. Willinger. 2005. The "robust yet fragile" nature of the Internet. *Proc. Nat. Acad. Sci. USA* **102**(41) 14497–14502.
- Erdős, P., A. Renyi. 1959. On random graphs I. *Publ. Math. (Debrecen)* **9** 290–297.
- Euler, L. 1736. Solutio problematis ad geometriam situs pertinentis. *Commetarii Academiae Scientiarum Petropolitanae* **8** 128–140.
- Evans, J., E. Minieka. 1992. *Optimization Algorithms for Networks and Graphs*. Marcel Dekker, New York.
- Fabrikant, A., E. Koutsoupias, C. Papadimitriou. 2002. Heuristically optimized trade-offs: A new paradigm for power-laws in the Internet. *Proc. ICALP*, Málaga, Spain, 110–122.
- Florian, M., D. Hearn. 1995. Network equilibrium models and algorithms. M. O. Ball, T. L. Magnanti, C. L. Monma, G. L. Nemhauser, eds. *Network Routing*. Chapter 6 in *Handbooks in Operations Research and Management Science*, Vol. 8. Elsevier, Amsterdam, 485–550.
- Ford, L. R., D. R. Fulkerson. 1962. *Flows in Networks*. Princeton University Press, Princeton, NJ.
- Giannessi, F., A. Maugeri, eds. 1995. *Variational Inequalities and Network Equilibrium Problems*. Plenum Press, New York.
- Girvan, M., M. E. J. Newman. 2002. Community structure in social and biological networks. *Proc. Nat. Acad. Sci. USA* **99** 7821–7826.
- Gleiss, P. M., P. F. Stadler, A. Wagner, D. A. Fell. 2001. Relevant cycles in chemical reaction networks. *Adv. Complex Systems* **4** 207–226.
- Gray, E., J.-M. Seigneur, Y. Chen, C. Jensen. 2003. Trust propagation in small worlds. P. Nixon, S. Terzis, eds. *Proc. First Internat. Conf. Trust Management*, Heraklion, Crete, Greece, 239–254.
- Hitchcock, F. L. 1941. The distribution of a product from several sources to numerous facilities. *J. Math. Phys.* **20** 224–230.
- Iamnitchi, A., M. Ripeanu, I. Foster. 2004. Small-world file-sharing communities. *Proc. IEEE INFOCOM*, Hong Kong. Vol. 2, 952–963.
- Jackson, M. O. 2006. The economics of social networks. R. Blundell, W. Newey, T. Persson, eds. *Advances in Economics and Econometrics, Theory and Applications*. Cambridge University Press, Cambridge, UK, 1–56.
- Jin, S., A. Bestavros. 2002. Small-world Internet topologies: Possible causes and implications on scalability of end-system multicast. Technical report BUCS-2002-004, Boston University, Boston.
- Johari, R., S. Mannor, J. N. Tsitsiklis. 2005. Efficiency loss in a network resource allocation game: The case of elastic supply. *IEEE Trans. Auto. Control* **50**(11) 1712–1724.
- Jovanović, M. A., F. S. Annexstein, K. A. Berman. 2001. Modeling peer-to-peer network topologies through "small-world" models and power laws. Presented at *IX Telecomm. Forum TELFOR 2001*, Belgrade, Serbia.
- Kantorovich, L. V. 1939. Mathematical methods in the organization and planning of production. Publication House of the Leningrad University. Translated in *Management Sci.* 1960 **6**(4) 366–422.
- Kauffman, S. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.
- Keller, E. F. 2005. Revisiting "scale-free" networks. *BioEssays* **27**(10) 1060–1068.
- Kelly, F. P. 2001. Mathematical modelling of the Internet. B. Engquist, W. Schmid, eds. *Mathematics Unlimited—2001 and Beyond*. Springer-Verlag, Berlin, 685–702.
- Kelly, F. P., A. Maulloo, D. Tan. 1998. Rate control in communication networks: Shadow prices, proportional fairness and stability. *J. Oper. Res. Soc.* **49** 237–252.
- Klemm, K., V. M. Eguiluz. 2002. Highly clustered scale-free networks. *Phys. Rev. E* **65** 036123.
- Koopmans, T. C. 1947. Optimum utilization of the transportation system. *Proc. Internat. Statist. Conf.*, Washington, D.C. Also in *Econometrica* 1949 **17** 136–145.
- Li, L., D. L. Alderson, J. C. Doyle, W. Willinger. 2006. Towards a theory of scale-free graphs: Definition, properties, and implications. *Internet Math.* **2**(4) 431–523.
- Li, L., D. Alderson, W. Willinger, J. Doyle. 2004. A first principles approach to understanding router-level topology. *ACM Comp. Comm. Rev.* **34** 3–14.
- Liljeros, F., C. R. Edling, L. A. N. Amaral, H. E. Stanley. 2001. The web of human sexual contacts. *Nature* **411** 907–908.

- Low, S. H., R. Srikant. 2004. A mathematical framework for designing a low-loss, low-delay Internet. *Networks and Spatial Econom.* **4** 75–101.
- Mandelbrot, B. 1997. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. Springer-Verlag, New York.
- Mitzenmacher, M. 2004. A brief history of generative models for power law and lognormal distributions. *Internet Math.* **1**(2) 226–251.
- Mitzenmacher, M. 2006. The future of power-law research. *Internet Math.* **2**(4) 525–534.
- Nagurney, A. 1993. *Network Economics: A Variational Inequality Approach, Advances in Computational Economics*, Vol. 1. Kluwer Academic Publishers, Boston. Second revised edition Vol. 10, Dordrecht, The Netherlands, 1999.
- Nagurney, A. 2003. Financial and economic networks: An overview. *Innovations in Financial and Economic Networks. New Dimensions in Networks*. Edward Elgar Publishing Inc., Cheltenham, UK, 1–26.
- Nagurney, A. 2006. *Supply Chain Network Economics: Dynamics of Prices, Flows, and Profits*. Edward Elgar Publishing Inc., Cheltenham, UK.
- National Research Council, Committee on Network Science for Future Army Applications. 2006. *Network Science*. The National Academies Press, Washington, D.C.
- Newman, M., A.-L. Barabási, D. J. Watts, eds. 2006. *The Structure and Dynamics of Networks*. Princeton University Press, Princeton, NJ.
- Newman, M. E. J. 2001. The structure of scientific collaboration networks. *Proc. Nat. Acad. Sci. USA* **98** 404–409.
- Newman, M. E. J. 2003. The structure and function of complex networks. *SIAM Rev.* **45** 167–256.
- Newman, M. E. J. 2005. Power laws, Pareto distributions and Zipf's law. *Contemporary Phys.* **46** 323–351.
- Newman, M. E. J., D. J. Watts. 1999a. Renormalization group analysis of the small-world network model. *Phys. Lett. A* **263** 341–346.
- Newman, M. E. J., D. J. Watts. 1999b. Scaling and percolation in the small-world network model. *Phys. Rev. E* **60** 7332–7342.
- Newman, M. E. J., C. Moore, D. J. Watts. 2000. Mean-field solution of the small-world network model. *Phys. Rev. Lett.* **84** 3201–3204.
- Newman, M. E. J., D. J. Watts, S. H. Strogatz. 2002. Random graph models of social networks. *Proc. Nat. Acad. Sci. USA* **99** 2566–2572.
- Ottino, J. 2004. Engineering complex systems. *Nature* **427**(29) 399.
- Pastor-Satorras, R., A. Vespignani. 2001. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.* **86**(14) 3200–3203.
- Pastor-Satorras, R., A. Vespignani. 2002. Epidemics and immunization in scale-free networks. S. Bornholdt, H. G. Schuster, eds. *Handbook of Graphs and Networks: From the Genome to the Internet*. Wiley-VCH, Weinheim, Germany, 111–130.
- Pastor-Satorras, R., A. Vespignani. 2004. *Evolution and Structure of the Internet: A Statistical Physics Approach*. Cambridge University Press, Cambridge, UK.
- Quesnay, F. 1758. *Tableau Economique*. Reproduced in facsimile with introduction by H. Higgs, British Economic Society, 1895, London.
- Ran, B., D. E. Boyce. 1996. *Modeling Dynamic Transportation Networks*. Springer-Verlag, Berlin.
- Roughgarden, T. 2005. *Selfish Routing and the Price of Anarchy*. MIT Press, Boston.
- Schrijver, A. 2002. On the history of the transportation and maximum flow problems. *Math. Programming Ser. B* **91** 437–445.
- Solé, R., S. Valverde. 2001. Information transfer and phase transitions in a model of Internet traffic. *Physica A* **289** 595–695.
- Stauffer, D., A. Aharony. 1992. *Introduction to Percolation Theory*, 2nd ed. Taylor and Francis, London.
- Tanaka, R. 2005. Scale-rich metabolic networks. *Phys. Rev. Lett.* **94** 168101.
- Tolstoj, A. N. 1930. Methods of finding the minimal total kilometrage in cargo transportation planning in space. *TransPress of the National Commissariat of Transportation* (Russian). Moscow, 23–55.
- Veres, A., M. Boda. 2000. The chaotic nature of TCP congestion control. *Proc. IEEE INFOCOM*, Tel Aviv, Israel.
- Wagner, A., D. Fell. 2001. The small world inside large metabolic networks. *Proc. Roy. Soc. London Ser. B* **268** 1803–1810.
- Watts, D. 1999a. Personal correspondence.
- Watts, D. J. 1999b. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton University Press, Princeton, NJ.
- Watts, D. J. 2003. *Six Degrees: The Science of a Connected Age*. W. W. Norton & Company, New York.
- Watts, D. J. 2004. The "new" science of networks. *Annual Rev. Sociol.* **30** 243–270.
- Watts, D. J., S. H. Strogatz. 1998. Collective dynamics of "small world" networks. *Nature* **393** 440–442.
- Waxman, B. M. 1988. Routing of multipoint connections. *IEEE J. Selected Areas in Comm.* **6**(9) 1617–1622.
- Weaver, W. 1948. Science and complexity. *Amer. Scientist* **36** 536–544.
- Wei, D. X., C. Jin, S. H. Low, S. Hegde. 2006. FAST TCP: Motivation, architecture, algorithms, performance. *IEEE/ACM Trans. Networking* **14**(6) 1246–1259.
- Willinger, W., D. Alderson, L. Li, J. C. Doyle. 2004. More "normal" than normal: Scaling distributions and complex systems. R. G. Ingalls, M. D. Rossetti, J. S. Smith, B. A. Peters, eds. *Proc. 2004 Winter Sim. Conf.*, Washington, D. C., 130–141.
- Willinger, W., R. Govindan, S. Jamin, V. Paxson, S. Shenker. 2002. Scaling phenomena in the Internet: Critically examining criticality. *Proc. Nat. Acad. Sci. USA* **99**(Suppl. 1) 2573–2580.