Topologies of Complex Networks: Functions and Structures

Thesis by

Lun Li

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

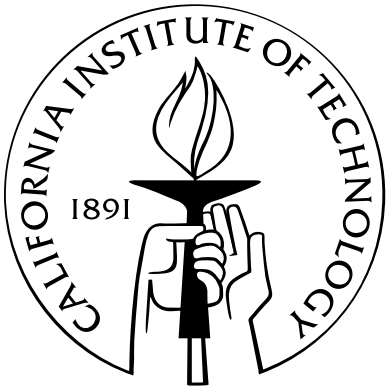

California Institute of Technology

Pasadena, California

2007

(Defended April 16, 2007)

To my parents Guobin and Guoqing,

my husband Zhenyu,

and my daughter Elaine.

Acknowledgements

I would like to thank my advisor Professor John C. Doyle and coadvisor Professor Steven H. Low from the bottom of my heart. Without their guidance, support and encouragement, this thesis would never have come true. John has provided me intellectual freedom, deep insights, and great pictures in my research. His endless effort in pursuing truth, devotion to athletics, and attention to charity has always been a source of inspiration. Steven has helped me open a door to study theoretical technical details from problem formulations to rigorous proofs. I have been fortunate to have both of them as my advisors.

Special thanks go to Dr. David Alderson and Dr. Walter Willinger, for enjoyable and fruitful collaborations that contributed to a major part of this thesis. My gratitude also extends to Professor Richard Murray, Professor Tracey Ho and Dr. Walter Willinger for serving on my thesis committee.

I want to thank my advisor Professor Andy Packard at U.C. Berkeley for guiding my first years in the United States, and thank my mentors and co-workers at Microsoft Research for the summer of 2006, specially to Ranveer Chandra, Srikanth Kandula and Ming Zhang.

Thanks also go to Gloria Bain, Christine Ortega, Betta Dawson, Peggy, Charmaine Boyd, Linda Dozsa for helping me with all sorts of problems. I would also like to thank my friends and colleagues in Netlab and CDS, to Jiantao Wang for the collaboration of the cross-layer optimization project and Antonis Papachristodoulou for the congestion control work, to Xin Liu, Sarah Tan, Qian Li, Qiang Yang, Xiaolan Wu, Huirong Ai, Hualin Ye, Rui Zhang-Shen, Qian Zhao, Jiantao Wang, Cheng Jin, Kevin Tang, David Wei, Chang Liu, Yindi Jin, Stephen Prajna, Aristotelis Asimakopoulos, Lachlan Andrew, Morr Mehyar, Zhipu Jin, Lijun Chen, Maryam Fazel, Denice Gayme, Ling Shi, Zhengrong Wang, Xiaoli Feng, Yan Wu and Xiang Li for the friendship and support.

Last but not least, I thank my parents for their help and care, my dear husband for his love and support, and my daughter for my source of happiness. This thesis is dedicated to them.

Abstract

During the last decade, significant efforts have been made toward improving our understanding of the topological structures underlying complex networks and illuminating some of the intriguing large-scale properties exhibited by these systems. The dominant theme of these efforts has been on studying the graph-theoretic properties of the corresponding connectivity structures and on developing universal theories and models that transcend system-specific details and describe the different systems well in a statistical sense.

However, in this thesis we argue that these efforts have had limited success and are in need of substantial correction. Using a highly engineered system, the Internet, as a case study we demonstrate that networks are designed for a purpose, and ignoring that aspect or obscuring it with the use of some generic but random mechanism can result in models that misrepresent what matters for system functions. By accounting in a minimal manner for both the functional requirements and structural features inherent in the design of an engineered system, we propose an alternative, optimization-based modeling approach that highlights the necessary trade-offs between system performance and the technological and economic constraints that are crucial when designing the system. We show that our proposed approach yields network models that not only match the large-scale graph-theoretic properties of measured router-level topologies well but are also fully consistent with engineering intuition and networking reality, especially as far as their performance aspects and robustness properties are concerned. In fact, we show that our design-inspired network models can be easily distinguished from previously considered probabilistic network models and efficiently achieve the level of performance for which they were designed in the first place.

While this thesis focuses on the Internet, it has much broader implications for com-

plex networks and graph theory generally. To better differentiate between different graphs that are identical in certain graph statistics, we introduce a structural metric, the *s*-metric, and demonstrate that it provides insights into the diversity of graphs constrained by certain common properties and sheds new light on many classic graph concepts such as the various notions of self-similarity, likelihood, and assortativity. Our *s*-metric clarifies much of the confusion surrounding the sensational qualitative claims in the current graph theory literature for complex networks and offers a rigorous and quantitative alternative.

Moreover, to examine the space of graphs that satisfy certain common properties, we propose a new approach that is based on establishing a link between two graphs if and only if one can be obtained from the other via a local transformation. Exploring the resulting connected space of graphs by dividing it into countable subspaces provides a much clearer picture on the whole space. We also show that this space of graphs has a rich and interesting structure and that some properties of the latter can be related to features of the individual graphs in this space (e.g., degree variability of a node g in the space of graphs and the *s*-metric for g).

Contents

Acknowledgements **v**

Abstract **vii**

1 Introduction **1**

1.1 Introduction to Complex Networks 1

1.2 Summary of Main Results 4

1.3 Thesis Organization 7

2 Background and Related Work **9**

2.1 Power-law and Scaling Behavior 9

2.1.1 High Variability 12

2.1.2 Cumulative vs. Noncumulative Log-log Plots 13

2.1.3 More “Normal” than Normal 16

2.1.4 Scaling Degree Sequence and Degree Distribution 18

2.2 Scale-Free Networks 20

2.2.1 Main Scale-Free Properties 21

2.2.2 Existing Scale-Free Literature 21

2.3 Summary 26

3 Internet Router-Level Topology **27**

3.1 Introduction 28

3.2 Previous Work on Internet Topology 30

3.3 A First-Principles Approach 33

3.3.1 Technology Constraints . . . 33
3.3.2 Economic Considerations . . . 36
3.3.3 Service Requirements . . . 37
3.3.4 Heuristically Optimal Networks . . . 39
3.4 Evidence from Real Internet . . . 40
3.4.1 The Abilene Network . . . 40
3.4.2 The CENIC Network . . . 41
3.4.3 A Tier-1 ISP Network . . . 43
3.5 Topology Metrics . . . 45
3.5.1 Commonly-Used Metrics . . . 45
3.5.2 Function-Related Metrics . . . 47
3.6 Comparing Topologies . . . 48
3.6.1 A First Example . . . 49
3.6.2 A Second Example . . . 54
3.7 Summary . . . 55

4 The Structural Metric **57**
4.1 Definitions and Properties . . . 58
4.1.1 Performance vs. Structural Metrics . . . 59
4.1.2 Joint Degree Distribution . . . 61
4.1.3 The $\mathbf{s}_{max}$ and $\mathbf{s}_{min}$ Graphs . . . 63
4.1.4 Diversity of Degree Sequence . . . 65
4.2 The **s**-metric and Self-Similarity . . . 69
4.2.1 Graph Trimming by Link Removal . . . 69
4.2.2 Coarse Graining by Collapsing Nodes . . . 70
4.2.3 Subgraph-Based Motifs . . . 71
4.2.4 Degree-Preserving Rewiring . . . 74
4.3 The **s**-metric and Likelihood . . . 76
4.3.1 Probabilistic Approach . . . 76
4.3.2 Highly Likely Constructions . . . 80

4.4 The **s**-metric and Assortativity 82
4.4.1 Assortativity Definition 82
4.4.2 Measuring Against Background Sets 83
4.4.3 Empirical Results 86
4.5 Summary 90

5 The GRAPH of graphs **91**
5.1 Introduction 92
5.2 The General Flip 94
5.2.1 Transformation Standards 95
5.2.2 The General Flip 95
5.3 G1 101
5.3.1 Properties of G1 101
5.3.2 The Canonical Graph in G1 104
5.3.3 Exploring G1 107
5.4 The Degree Preserving Flip and G2 112
5.4.1 The Degree Preserving Flip 112
5.4.2 Properties of G2 115
5.5 Summary 117

6 Other Projects **119**
6.1 Cross-Layer Optimization in TCP/IP Networks 119
6.2 Methodological Frameworks For Internet Congestion Control 122
6.3 Detect User Perceived Faults Using Packet Traces 123

7 Conclusions and Future Work **127**
7.1 Conclusions 127
7.2 Future Directions 129
7.2.1 Internet Topology Generator 129
7.2.2 Apply the HOT Idea to Other networks 130
7.2.3 Benchmark Graphic Metrics 131

7.2.4 the Graph Space: Properties and Dynamics 132

A Constructing an s_{max} Graph **133**

A.1 Among "Unconstrained" Graphs 133

A.2 Among Graphs in $G(D)$ 134

A.2.1 The Algorithm 136

A.2.2 Among Connected, Acyclic Graphs 141

A.3 When $r_{min} = -1$ 142

Bibliography **144**

List of Figures

2.1 Plots of exponential y^e (black circles) and scaling y^s (blue squares) sequences. (a) Doubly logarithmic size-rank plot. (b) Semi-logarithmic size-rank plot. (c) Doubly logarithmic size-frequency plot. (d) Semi-logarithmic size-frequency plot. 14

3.1 Technology constraint for Cisco 12416 Gigabit Switch Router(GSR): degree vs. bandwidth. 34

3.2 Aggregate picture of router technology constraints. 36

3.3 Aggregate picture of end user connection bandwidths for the Internet. 37

3.4 Abilene network. Each node represents a router, and each link represents a physical connection between Abilene and another network. 40

3.5 CENIC backbone. 42

3.6 Configuration of CENIC routers. 43

3.7 Configuration of a Tier-1 commercial ISP. 44

3.8 Five networks having the same node degree distribution as shown in sub-figure (a) 50

3.9 (a) Distribution of end user bandwidths for all the five different networks; (b)-(f) Router utilization for each network. The colorscale of a router on each plot differentiates its bandwidth which is consistent with the routers in figure 3.8. 51

3.10 Robustness of network performance to router loss for the PA and HOT networks from Figure 3.8. 54

3.11 Distribution of node degree and end-user bandwidths for several topologies having the same core structure: (a) uniformly high bandwidth end users, (b) highly variable bandwidth end users, (c) uniformly low bandwidth end users. 55

4.1 Performance vs. the structural metric for each topology, plus other networks having the same node degree distribution obtained by pairwise random rewiring of links. ($S(g) = (s(g) - s_{\min}) / (s_{\max} - s_{\min})$) 60

4.2 Graph diversity and degree sequence diversity. (a) The $s_{\min}$ and $s_{\max}$ values in both $\mathcal{G}(D)$ and $G(D)$ vs. the $CV(D)$ of the corresponding degree sequence. Note that $s_{\min}^{G(D)} \approx s_{\min}^{\mathcal{G}(D)}$. (b) The $s_{\min}$ and $s_{\max}$ in $G(D)$, each normalized by their respective $s_{\max}$. 68

4.3 Three possible subgraph-based motifs in degree-preserving rewiring in acyclic graphs. 72

4.4 Results from Monte Carlo generation of preferential attachment graphs having 1000 nodes. Both the CDF and CCDF are shown. 81

4.5 Comparison between s and r with respect to CV 87

4.6 Graph diversity among nontrees. In this experiment, an additional $k(n-1)$ links were added to initial trees of size $n = 100$. (a) $k = 1$, $\langle d \rangle = 3.96$, $C_V^{\max} = 3.4451$. (b) $k = 2$, $\langle d \rangle = 5.94$, $C_V^{\max} = 2.7672$. (c) $k = 4$, $\langle d \rangle = 9.9$, $C_V^{\max} = 2.0701$. In the bottom graphs, variation is measured with $CV(D)$ while in the top graphs it is represented as the normalized $CV(D)/C_V^{\max}(D)$. . 89

5.1 The general flip. 96

5.2 A series of general flips to make node v connect to u_1. 98

5.3 A series of general flips to make any edge node v with degrees greater than one to connect to the second star node u_2. 99

5.4 The balancing process to add link (v_1, u_2) and remove link (v_2, v_3) 99

5.5 Graph probability vs. degree sequence variability. 111

5.6 The degree-preserving flip. 113

5.7 The bow-tie switch. 113

5.8 The degree-preserving flip for bow-tie switch when extra node is connected to the edge nodes of bow-tie (case 1). 114
5.9 The degree-preserving flip for bow-tie switch when extra nodes are connected to the center node (case 2). 115

List of Tables

4.1 The numbers of the three motifs and successively the number for each possible rewiring outcome in figure 4.3. 72

Chapter 1

Introduction

1.1 Introduction to Complex Networks

A network is a collection of network components representing its fundamental units and a set of connections characterizing any relationship between these components. Networks are ubiquitous, ranging from biological networks to social networks to technological networks. Examples of biological networks include the cellular network, which is an ensemble of genes, proteins and other molecules, and their interactions to regulate cell activities; a biological neural network consisting of functionally related neurons that perform a specific physiological function. The famous Erdös numbers in fact describe a social network where mathematicians are assigned numbers indicating the "collaboration distance" to a well-known mathematician Paul Erdös who wrote about 1500 papers in his life, mostly coauthored with others. As an example of a technological network, the Internet, one of the largest man-made networks, can be defined as a huge collection of millions of computers and routers connected by physical links, or in a more coarse-grain level, can be considered as consisting of thousands of administrative domains among which data are transferred. All these are just a small set of complex networks.

Over the last decade there has been significant interest and attention devoted toward understanding the infrastructure underlying complex networks, particularly their topologies and the large-scale properties that can be derived. The topology of a complex network is usually pictured as a graph, where nodes (or vertices) represent basic network components, and links (or edges) portray their interactions. Studying the topological structure of com-

plex networks has been one of the most fundamental steps to gain a basic understanding of certain aspects of real-world phenomena of many kinds due to a simple reason: network structures always affect network functions. For example, food webs are investigated as the road-maps through Darwin's entangled bank [90, 107]. Social network topologies can help to prevent pandemic influenza from spreading when available to health care. Studying the topological structure of complex networks also plays an important role to evaluate and design network regulations and protocols that run on top of them. Although the topology should not affect their correctness, it always affect their performance. Understanding topology of complex networks can also protect networks from failures and attacks, so as to achieve a better design and evolution of networks.

However, studying topologies of complex networks has proved to be a challenging problem. Since a large-scale network is usually a collection of thousands or millions of nodes, there is no single place from which one can obtain a complete picture of the topology. Moreover, networks are dramatically changing and constantly evolving. For example a web page on the World Wide Web can be created or removed on a daily basis, and it is impossible to obtain a snapshot of this network. Furthermore, because the network does not lend itself naturally to direct inspection, the task of "discovering" topologies has been left to experimentalists who develop more or less sophisticated methods to infer this topology from appropriate network measurements. Because of the elaborate nature of the network, there are a multitude of possible measurements that can be made, each having its own strengths, weaknesses, and idiosyncrasies, and each resulting in a distinct view of the network topology.

Due to these challenges, the recent use of network models to describe complex systems has emphasized the study of graph theoretic properties as a means to characterize the similarities and differences in the structures and the functions of systems across a variety of domains [42, 95, 10, 96, 43, 27, 103]. Considerable effort has been directed at the empirical analysis of graph theoretic properties of real systems and at trying to find unifying properties across many complex networks. Even more attention has focus on to develop generic and universal models to attempt to explain such unifying properties, so as to infer more properties that are not easy to obtain by empirical analysis. An implicit assumption

in many of these works is that graph theoretic properties adequately capture key system features in order to serve as a basis for comparison and contrast.

One of the most celebrated properties discovered across many topologies of complex networks is the high variability in degree distributions, and this high variability significantly deviates from the low variability distribution such as Poisson distribution and exponential distribution in traditional random networks. In particular, these highly variable distributions follow a power-law relationship in many networks, such as both the router-level and AS-level topologies of the Internet [48], the World Wide Web [11], the network of citation between scientific papers [108], metabolic reaction network [59], and the telephone call graph [7].

Since traditional graph theory on regular graph or random graph cannot explain the high variability of degree sequence, the discovery of the power-law degree distribution has stimulated a great deal of work in the construction of the so-called "scale-free" networks, aiming to match the power-law distribution and other large scale statistical properties, as well as to provide a universal theory to understand all complex networks. The most famous model was proposed by Barabási and Albert [20], who describe a growing process called preferential attachment for a complex network in which a new node is added to the network with probability proportional to the degree of existing nodes. As the high degree nodes can connect to more and more nodes, i.e. rich and richer, these nodes significantly contribute to the high variability in the power-law distribution. Since then, numerous refinements and modifications to the original Barabási-Albert construction have been proposed and have resulted in many types of scale-free network models that can reproduce power-law degree distributions with different variation, for example the ability to tune the parameters of the power-law distribution, in order to agree with different complex networks [10].

Despite of these variations, scale-free networks share many common features: the most attractive one is that it has power-law degree distribution which makes it a plausible model for many complex networks. In fact, scale-free theory has dominated the current literature of complex networks and has been considered as the universal law for any large-scale networks since none of the previous graph theory can explain the power-law degree distribution. Moreover, scale-free graphs are claimed to exhibit a host of startling "emergent"

consequences of universal relevance, including intriguing self-similar and fractal properties, small-world characteristics [15], and "hublike" cores. Perhaps the central claim for scale-free graphs is that they have highly connected hubs, which "hold the network together." [11] As noted, the structures of such networks are highly vulnerable (i.e., can be fragmented) to attacks that target these hubs [11]. At the same time, they are resilient to attacks that knock out nodes at random, since a randomly chosen node is unlikely to be a hub, and thus its removal has minimal effect on network connectivity. In the context of the Internet, where scale-free graphs have been proposed as models of the router-level topology [125], this has been touted as "the Achilles' heel of the Internet" [11], a vulnerability that has presumably been overlooked by networking engineers. Proponents of this modeling framework have further suggested that the emergent properties of scale-free graphs contributes to truly universal behaviors in complex networks [25] and that preferential attachment as well is a universal mechanism at work in the evolution of these networks [62, 43].

Notwithstanding the potential pitfalls of reducing a complex system (e.g., one that may involve heterogeneous components, layered architectures, and feedback dynamics) to a simple graph [44, 116, 71], there exists the practical problem that many descriptions based on aggregate statistics do not uniquely characterize the system of interest. In fact, there often exists considerable diversity among graphs that share any single statistical feature, particularly when viewed through the lens of a specific application domain.

1.2 Summary of Main Results

In this thesis, we show that the current models and theories for understanding complex networks are incomplete and in need for substantial corrective actions, for both functional and structural reasons.

Using the Internet, a highly engineered network topology as an example, we illustrate the role of network functions in modeling network topologies. Highly engineered networks are designed for a purpose, and ignoring that aspect or obscuring it with the help of some generic but random mechanism can result in models that are meaningless from an engi-

neering perspective. By leveraging minimal functional requirements and constraints faced by network engineers when designing the current Internet, we propose a *Heuristic Optimal Topology* (HOT) model which considers the trade-off between performance optimization and technology and economic constraint. The HOT model captures all the important large-scale graphic properties as previous models, yet has fundamental functional and structural differences. In contrast to the highly connected hubs at the center in scale-free networks, all the high degree nodes are at the edges to aggregate as many end users as possible, while the core consists meshlike low degree nodes to carry as much traffic as possible. The advantages in performance and robustness of the HOT model are essential to the Internet design and consistent with engineering reality, while scale-free models have a such bad performance as to lend them no reason to exist in real Internet.

Our study shows there is enough diversity among graphs having the same power-law node degree distribution that, although indistinguishable when viewed by this aggregate statistic, these graphs can actually be interpreted as "opposites" when viewed from an engineering perspective that incorporates technology constraints and is motivated by throughput performance. We further introduce a structural metric (the *s*-metric) to quantitatively characterize the extend to which a network has highly connected hubs, i.e., a graph is scale-free. The *s*-metric allows us to differentiate between all simple and connected graphs constrained by common macroscopic connectivity, which is of particular interest when the graphs satisfy highly variable degree sequence. We show that the *s*-metric, and in particular an $s_{\max}$ graph (a graph with maximal *s*-value, therefore a perfect scale-free graph), is relevant for many commonly studied graph properties. First, high degree nodes in the $s_{\max}$ graph have high *centrality*, and for trees this relationship was shown to be monotonic (see [70]). Second, $s_{\max}$ graphs are *self-similar* under appropriately defined operations of trimming, coarse graining, network motifs and random rewiring. Finally, the $s_{\max}$ graph has the highest likelihood of being generated by the most popular power-law degree-based graph model. As we aware, the *s*-metric is the first metric that is introduced targeting to the background set with all simple connected graphs with common high variable degree sequence. In the process of investigating the relationship between the *s*-value and the graph assortativity, we discover that the assortativity, as a popular metric directly borrowed from classic

graph theory where all the graphs have low variability, is extremely misleading when it is applied to high variability case. Therefore we advocate the importance of choosing an appropriate "background set" when evaluating a graph, as well as the importance of making sure that the comparative analysis of two graphs is conducted with respect to an appropriate reference. In this regard, not all graph theoretic measures have an obvious interpretation or are directly comparable.

We propose performance-related metrics for the Internet router-level topology to characterize the function of this specific complex network, and the *s*-metric as an orthogonal view to depict the structural differences of all the networks with the same high variability degree distributions. These macroscopic properties are fundamental to understand the functions and structures of complex networks, and together project the extremely diverse space of graphs into a two-dimensional plane. As a complement of studying the topology of complex networks, we further explore the space of graphs from a microscopic view, where all the graphs are connected according to their structural relationship, defined by some local transformations from one graph to another. This connected space of graphs is noted as a *GRAPH* of graphs in which each node is a graph and each link represents a transformation between the corresponding two graphs. Although the GRAPH of graphs is much more complicated than each individual graph, we can break this giant space into many countable subspaces having common properties, and enumerate all the possible graphs in each subspace. As a result, our work depicts a much more detailed and cleared picture of the space of graphs which has never been explored carefully before, and lends perspective on the structural relationship among all the graphs in the same domain space. Interestingly enough, we find that many properties of the GRAPH of graphs have direct connections to the properties of graphs inside it. For example, when the GRAPH of graphs represents a domain space of all the simple and connected graphs with the same numbers of nodes and links, the degree of a node in the GRAPH is most relevant to the degree variability of the graph that node represents. While further constrained to the graphs with the same degree sequence, many properties of the GRAPH are related to the graph *s*-metric.

1.3 Thesis Organization

The rest of thesis is organized as follows:

We provides basic background knowledge to study complex networks in chapter 2. Specifically, we give a precise definition of power-law distribution and its important properties. Then we overview scale-free networks literature, and the important properties of scale-free networks.

In chapter 3, we emphasize the functions and constraints for router-level topology in Internet and propose an optimization-based model which combines objectives, constraints and other drivers of engineering design. We compare our model and scale-free model by evaluating their performance related metrics such as total throughput and robustness to worst case attack. We also explore several real Internet topologies to show their consistence with our model.

In chapter 4, we propose a structural metric, the *s*-metric to differentiate graphs with the same degree sequence. We investigate the relationship between *s*-metric and scale-free network and show that *s*-metric is in fact a measure of the extend to which a graph is scale-free. We further explore the detailed relationship between *s*-metric and some well-known graph properties like similarity, likelihood and assortativity.

We present the GRAPH of graphs and the way to construct it in chapter 5. We describe a method to explicitly calculate the probability of graphs with different degree variabilities, and provide a much clearer picture of the space of graphs with the same numbers of nodes and links. We discuss properties of the GRAPH and their implications to variability of degree and the *s*-metric of a graph.

In chapter 6, several of my other projects are described briefly. If the topologies of complex networks can be thought as my horizontal research, my vertical interests span several layers of Internet, including the application layer failure detection and identification, transport layer congestion control, cross-layer joint optimization for TCP/IP.

We conclude and point out several related future directions at the end.

Chapter 2

Background and Related Work

This chapter provides the necessary background for our investigation of complex networks. In particular, we present some basic definitions and results regarding the power-law distribution (also called Scaling distribution) and scale-free network models.

We point out the scaling and highly variable behaviors of the power-law distributions, and comment on some common mistakes for plotting power-law relations in a log-log scale. We also talk about "more normal than normal," the reason why power-law distribution is ubiquitous in large scale networks from a pure mathematical point of view.

In the second part, we first describe the basic properties and claims of scale-free networks, and then review the existing scale-free network literature and present some of the most popular models. This is followed by a brief critique of the existing theory of scale-free networks in general.

2.1 Power-law and Scaling Behavior

A finite *sequence* $y = (y_1, y_2, \ldots, y_n)$ of real numbers, assumed without loss of generality always to be ordered such that $y_1 \geq y_2 \geq \ldots \geq y_n$, is said to follow a *power-law* or *scaling relationship* if

$$k = c{y_k}^{-\alpha}, \tag{2.1}$$

where k is (by definition) the *rank* of y_k, c is a fixed constant, and α is called the *scaling index*. Since $\log k = \log(c) - \alpha \log(y_k)$, the relationship for the rank k vs. y appears as a line of slope $-\alpha$ when plotted on a log-log scale. In this thesis, we refer to the relationship (2.1) as the *size-rank* (or *cumulative*) form of scaling. While the definition of scaling in (2.1) is fundamental to the exposition of our work, a more common usage of power-laws and scaling occurs in the context of random variables and their distributions. That is, assuming an underlying probability model P for a nonnegative random variable X, let $F(x) = P[X \leq x]$ for $x \geq 0$ denote the *(cumulative) distribution function (CDF) of* X, and let $\bar{F}(x) = 1 - F(x)$ denote the *complementary CDF (CCDF)*.

In this stochastic context, a random variable X or its corresponding distribution function F is said to follow a *power-law* or is *scaling* with index $\alpha > 0$ if, as $x \to \infty$,

$$P[X > x] = 1 - F(x) \approx cx^{-\alpha}, \tag{2.2}$$

for some constant $0 < c < \infty$ and a *tail index* $\alpha > 0$. Here, we write $f(x) \approx g(x)$ as $x \to \infty$ if $f(x)/g(x) \to 1$ as $x \to \infty$. For $1 < \alpha < 2$, F has infinite variance but finite mean, and for $0 < \alpha \leq 1$, F has not only infinite variance but also infinite mean. In general, all moments of F of order $\beta \geq \alpha$ are infinite. Since relationship (2.2) implies $\log(P[X > x]) \approx \log(c) - \alpha \log(x)$, doubly logarithmic plots of x vs. $1 - F(x)$ yield straight lines of slope $-\alpha$, at least for large x. In contrast, *exponential distributions* (i.e., $P[X > x] = e^{-\lambda x}$) result in approximately straight lines on semi-logarithmic plots.

If the derivative of the cumulative distribution function $F(x)$ exists, then $f(x) = \frac{d}{dx}F(x)$ is called the *(probability) density function* of X and implies that the stochastic cumulative form of scaling or size-rank relationship (2.2) has an equivalent *noncumulative* or *size-frequency* counterpart given by

$$f(x) \approx cx^{-(1+\alpha)}, \tag{2.3}$$

which appears similarly as a line of slope $-(1+\alpha)$ on a log-log scale. However, as discussed in more detail in section 2.1.2 below, the use of this noncumulative form of scaling has

been a source of many common mistakes in the analysis and interpretation of actual data and should generally be avoided.

Power-law distributions are called scaling distributions because the sole response to conditioning is a change in scale; that is, if the random variable X satisfies relationship (2.2) and $x > w$, then the conditional distribution of X given that $X > w$ is given by

$$P[X > x|X > w] = \frac{P[X > x]}{P[X > w]} \approx c_1 x^{-\alpha},$$

where the constant c_1 is independent of x and is given by $c_1 = 1/w^{-\alpha}$. Thus, at least for large values of x, $P[X > x|X > w]$ is identical to the (unconditional) distribution $P[X > x]$, except for a change in scale. In contrast, the exponential distribution gives

$$P(X > x|X > w) = e^{-\lambda(x-w)},$$

that is, the conditional distribution is also identical to the (unconditional) distribution, except for a change of location rather than scale. Thus we prefer the term *scaling* to *power-law*, but will use them interchangeably, as is common.

It is important to emphasize again the differences between these alternative definitions of scaling. Relationship (2.1) is *nonstochastic*, in the sense that there is no assumption of an underlying probability space or distribution for the sequence y, and in what follows we will always use the term *sequence* to refer to such a nonstochastic object y, and accordingly we will use *nonstochastic* to mean simply the absence of an underlying probability model. In contrast, the definitions in (2.2) and (2.3) are *stochastic* and require an underlying probability model. Accordingly, when referring to a random variable X we will explicitly mean an ensemble of values or realizations sampled from a common distribution function F, as is common usage. We will often use the standard and trivial method of viewing a nonstochastic model as a stochastic one with a singular distribution.

These distinctions between stochastic and nonstochastic models will be important in our work. Our approach allows for but does not require stochastics. In contrast, the scale-free literature almost exclusively assumes some underlying stochastic models, so we will

focus some attention on stochastic assumptions. Exclusive focus on stochastic models is standard in statistical physics, even to the extent that the possibility of nonstochastic constructions and explanations is largely ignored. This seems to be the main motivation for viewing the Internet's router topology as a member of an ensemble of random networks, rather than an engineering system driven by economic and technological constraints plus some randomness, which might otherwise seem more natural. Indeed, in current literature "random" is typically used more narrowly than stochastic to mean, depending on the context, exponentially, Poisson, or uniformly distributed. Thus phrases like "scale-free vs. random" (the ambiguity in "scale-free" notwithstanding) are closer in meaning to "scaling vs. exponential," rather than "nonstochastic vs. stochastic."

2.1.1 High Variability

An important feature of sequences that follow the scaling relationship (2.1) is that they exhibit *high variability*, in the sense that deviations from the average value or (sample) mean can vary by orders of magnitude, making the average largely uninformative and not representative of the bulk of the values. To quantify the notion of *variability*, we use the standard measure of *(sample) coefficient of variation*, which for a given sequence $y = (y_1, y_2, \ldots, y_n)$ is defined as

$$CV(y) = \sigma(y)/\bar{y}, \tag{2.4}$$

where $\bar{y} = n^{-1} \sum_{k=1}^{n} y_k$ is the average size or (sample) mean of y and $\sigma(y) = (\sum_{k=1}^{n}(y_k - \bar{y})^2/(n-1))^{1/2}$ is the (sample) standard deviation, a commonly-used metric for measuring the deviations of y from its average $\bar{y}$. The presence of high variability in a sequence of values often contrasts greatly with the typical experience of many scientists who work with empirical data exhibiting *low variability*—that is, observations that tend to concentrate tightly around the (sample) mean and allow for only small to moderate deviations from this mean value.

A standard ensemble-based measure for quantifying the variability inherent in a random

variable X is the *(ensemble) coefficient of variation CV(X)* defined as

$$CV(X) = \sqrt{\mathrm{Var}(X)}/E(X), \tag{2.5}$$

where $E(X)$ and $Var(X)$ are the (ensemble) mean and (ensemble) variance of X, respectively. If $x = (x_1, x_2, \ldots, x_n)$ is a realization of an independent and identically distributed (iid) sample of size n taken from the common distribution F of X, it is easy to see that the quantity $CV(x)$ defined in (2.4) is simply an estimate of $CV(X)$. In particular, if X is scaling with $\alpha < 2$, then $CV(X) = \infty$, and estimates $CV(x)$ of $CV(X)$ diverge for large sample sizes. Thus, random variables having a scaling distribution are extreme in exhibiting high variability. However, scaling distributions are only a subset of a larger family of *heavy-tailed distributions* (see [123] and references therein) that exhibit high variability. It turns out that some of the most celebrated claims on complex networks have as a necessary condition only the presence of high variability and not necessarily strict scaling per se. The consequences of this observation are far reaching, especially because they shift the focus from scaling relationships, their tail indices, and their generating mechanisms to an emphasis on heavy-tailed distributions and identifying the main sources of "high variability."

2.1.2 Cumulative vs. Noncumulative Log-log Plots

While in principle there exists an unambiguous mathematical equivalence between distribution functions and their densities, as in (2.2) and (2.3), no such relationship can be assumed to hold in general when plotting sequences of real or integer numbers or measured data cumulatively and noncumulatively. Furthermore, there are good practical reasons to avoid noncumulative or size-frequency plots altogether (a sentiment echoed in [97]), even though they are often used exclusively in some communities. To illustrate the basic problem, we first consider two sequences, y^s and y^e, each of length 1000, where $y^s = (y^s_1, \ldots, y^s_{1000})$ is constructed so that its values all fall on a straight line when plotted on doubly logarithmic (i.e., log-log) scale, i.e., y^s is a scaling sequence. Similarly, the values of the sequence $y^e = (y^e_1, \ldots, y^e_{1000})$ are generated to fall on a straight line when plotted on semi-logarithmic (i.e., log-linear) scale therefore y^e is an exponential sequence. The full

sequences are plotted in figure 2.1. In particular, the doubly logarithmic plot in figure 2.1(a) shows the cumulative or size-rank relationships associated with the sequences y^s and y^e. In full agreement with the underlying generation mechanisms, plotting on doubly logarithmic scale the rank-ordered sequence of y^s vs. rank k results in a straight line; i.e., y^s is scaling (to within integer tolerances). The same plot for the rank-ordered sequence of y^e has a pronounced concave shape and decreases rapidly for large ranks—strong evidence for an exponential size-rank relationship. Indeed, as shown in figure 2.1(b), plotting on semi-logarithmic scale the rank-ordered sequence of y^e vs. rank k yields a straight line; i.e., y^e is exponential (to within integer tolerances). The same plot for y^s shows a pronounced convex shape and decreases very slowly for large rank values—fully consistent with a scaling size-rank relationship.

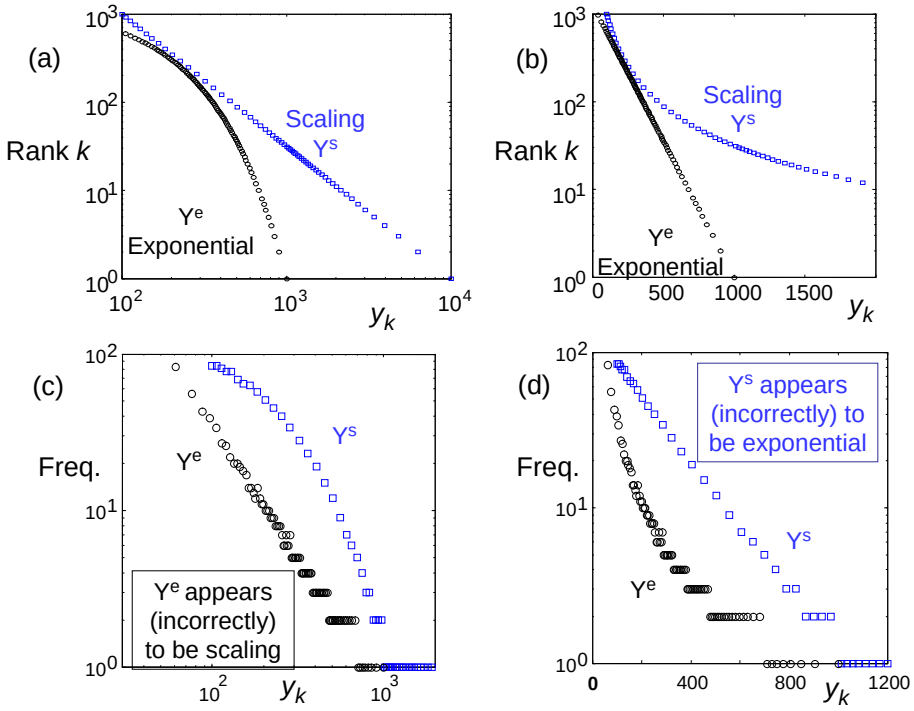

Figure 2.1: Plots of exponential y^e (black circles) and scaling y^s (blue squares) sequences. (a) Doubly logarithmic size-rank plot. (b) Semi-logarithmic size-rank plot. (c) Doubly logarithmic size-frequency plot. (d) Semi-logarithmic size-frequency plot.

To highlight the basic problem caused by the use of noncumulative or size-frequency relationships, consider Figure 2.1(c) and (d) that show on doubly logarithmic scale and

semi-logarithmic scale, respectively, the noncumulative or size-frequency plots associated with the sequences y^s and y^e: the largest value of y^s is plotted on the x-axis and has frequency 1 (y-axis), the second largest value of y^s has also frequency 1, etc., until the end where the smallest value of y^s happens to occur 84 times (to within integer tolerances). Similarly for y^e, the smallest value happens to occur 180 times. It is common to conclude incorrectly from plots such as these, for example, that the sequence y^e is scaling (i.e., plotting on doubly logarithmic scale size vs. frequency results in an approximate straight line) and the sequence y^s is exponential (i.e., plotting on semi-logarithmic scale size vs. frequency results in an approximate straight line)—exactly the opposite of what is correctly inferred about the sequences using the cumulative or size-rank plots in figure 2.1(a) and (b).

In contrast to the size-rank plots of the style in figure 2.1(a)-(b) that depict the raw data itself and are unambiguous, the use of size-frequency plots as in figure 2.1(c)-(d), while straightforward to describe low variability data, creates ambiguities and can easily lead to mistakes when applied to high variability data. First, for high precision measurements it is possible that each data value appears only once in a sample set, making raw frequency-based data rather uninformative. To overcome this problem, a typical approach is to group individual observations into one of a small number of *bins* and then plot for each bin (x-axis) the relative number of observations in that bin (y-axis). The problem is that choosing the size and boundary values for each bin is a process generally left up to the experimentalist, and this *binning process* can dramatically change the nature of the resulting size-frequency plots as well as their interpretation.

These examples have been artificially constructed specifically to dramatize the effects associated with the use of cumulative or size-rank vs. noncumulative or size-frequency plots for assessing the presence or absence of scaling in given sequence of observed values. While they may appear contrived, errors such as those illustrated in figure 2.1 are easy to make and are widespread in the complex systems literature. In fact, determining whether a realization of a sample of size n generated from one and the same (unknown) underlying distribution is consistent with a scaling distribution and then estimating the corresponding tail index α from the corresponding size-frequency plots of the data is even more unreliable.

Even under the most idealized circumstances using synthetically generated pseudorandom data, size-frequency plots can mislead as shown in the following easily reproduced numerical experiments (see [70] for details).

The log-log size-frequency plot Figure 2.1(c), however could be used incorrectly to claim that the data is consistent with a scaling distribution, a surprisingly common error in the scale-free and broader complex systems literature. Thus even if one a priori assumes a probabilistic framework, (cumulative) size-rank plots are essential for reliably inferring and subsequently studying high variability, and they therefore are used exclusively in this thesis.

2.1.3 More "Normal" than Normal

While power-laws in event size statistics in many complex interconnected systems have recently attracted a great deal of popular attention, some of the aspects of scaling distributions that are crucial and important for mathematicians and engineers have been largely ignored in the larger complex systems literature. This subsection will briefly review one aspect of scaling that is particularly revealing in this regard and is a summary of results described in more detail in [80, 123].

Gaussian distributions are universally viewed as "normal," mainly due to the well-known Central Limit Theorem (CLT). In particular, the ubiquity of Gaussians is largely attributed to the fact that they are invariant and attractors under aggregation of summands, required only to be independent and identically distributed (iid) and have finite variance [51]. Another convenient aspect of Gaussians is that they are completely specified by mean and variance, and the CLT justifies using these statistics whenever their estimates robustly converge, even when the data could not possibly be Gaussian. For example, much data can only take positive values (e.g., connectivity) or have hard upper bounds but can still be treated as Gaussian. It is understood that this approximation would need refinement if additional statistics or tail behaviors are of interest. Exponential distributions have their own set of invariance properties (e.g., conditional expectation) that make them attractive models in some cases. The ease by which Gaussian data is generated by a variety of mechanisms

means that the ability of any particular model to reproduce Gaussian data is not counted as evidence that the model represents or explains other processes that yield empirically observed Gaussian phenomena. However, a disconnect often occurs when data have high variability, that is, when variance or coefficient of variation estimates do not converge. In particular, the above type of reasoning is often misapplied to the explanation of data that are approximately scaling, for reasons that we will discuss below.

Much of science has focused so exclusively on low variability data and Gaussian or exponential models that low variability is not even seen as an assumption. Yet much real world data has extremely high variability as quantified, for example, via the coefficient of variation defined in (2.5). When exploring stochastic models of high variability data, the most relevant mathematical result is that the CLT has a generalization that relaxes the finite variance (e.g., finite *CV*) assumption, allows for high variability data arising from underlying infinite variance distributions, and yields *stable laws* in the limit. There is a rich and extensive theory on stable laws (see for example [109]), which we will not attempt to review, but mention only the most important features. Recall that a random variable U is said to have a *stable law (with index* $0 < \alpha \leq 2$*)* if for any $n \geq 2$, there is a real number d_n such that

$$U_1 + U_2 + \cdots + U_n \stackrel{d}{=} n^{1/\alpha} U + d_n,$$

where $U_1,\ U_2,\ \ldots,\ U_n$ are independent copies of U, and where $\stackrel{d}{=}$ denotes equality in distribution. Following [109], the stable laws on the real line can be represented as a four-parameter family $S_\alpha(\sigma, \beta, \mu)$, with the *index* α, $0 < \alpha \leq 2$; the *scale parameter* $\sigma > 0$; the *skewness parameter* β, $-1 \leq \beta \leq 1$; and the *location (shift) parameter* μ, $-\infty < \mu < \infty$. When $1 < \alpha < 2$, the shift parameter is the mean, but for $\alpha \leq 1$, the mean is infinite. There is an abrupt change in tail behavior of stable laws at the boundary $\alpha = 2$. While for $\alpha < 2$, all stable laws are scaling in the sense that they satisfy condition (2.2) and thus exhibit infinite variance or high variability; the case $\alpha = 2$ is special and represents a familiar, not scaling distribution—the Gaussian (normal) distribution, i.e., $S_2(\sigma, 0, \mu) = N(\mu, 2\sigma^2)$, corresponding to the finite variance or low variability case. While with the exception of Gaussian, Cauchy, and Levy distributions, the distributions of stable random variables are

not known in closed form, they are known to be the only fixed points of the renormalization group transformation and thus arise naturally in the limit of properly normalized sums of iid scaling random variables. From an unbiased mathematical view, the most salient features of scaling distributions are this and additional strong invariance properties (e.g., to marginalization, mixtures, maximization), and the ease with which scaling is generated by a variety of mechanisms [80, 123]. Combined with the abundant high variability in real world data, these features suggest that scaling distributions are in a sense more "normal" than Gaussians and that they are convenient and parsimonious models for high variability data in as strong a sense as Gaussians or exponentials are for low variability data.

While the ubiquity of scaling is increasingly recognized and even highlighted in the physics and the popular complexity literature [16, 31, 19, 18], the deeper mathematical connections and their rich history in other disciplines have been largely ignored, with serious consequences. Models of complexity using graphs, lattices, cellular automata, and sandpiles preferred in physics and the standard laboratory-scale experiments that inspired these models exhibit scaling only when finely tuned in some way. So even when accepted as ubiquitous, scaling is still treated as arcane and exotic, and "emergence" and "self-organization" are invoked to explain how this tuning might happen [14]. But given the strong invariance properties of scaling distributions, as well as the multitude of diverse mechanisms by which scaling can arise in the first place [97], it becomes clear that an ability to generate scaling distributions "explains" little, if anything. Once high variability appears in real data, then scaling relationships become a natural outcome of the processes that measure them.

2.1.4 Scaling Degree Sequence and Degree Distribution

Statistical features of graph structures that have received extensive treatment include the size of the largest connected component, link density, node degree relationships, the graph diameter, the characteristic path length, the clustering coefficient, and the betweenness centrality (for a review of these and other metrics see [10, 96, 43]). However, the single feature that has received the most attention is the distribution of node degree and whether

or not it follows a power-law.

For a graph with n nodes, let d_i denote the degree of node i, $1 \leq i \leq n$, and call $D = \{d_1, d_2, \ldots, d_n\}$ the *degree sequence* of the graph, again assumed without loss of generality always to be ordered $d_1 \geq d_2 \geq \ldots \geq d_n$. We will say a graph has *scaling degree sequence* D (or D *is scaling*) if for all $1 \leq k \leq n_s \leq n$, D satisfies a *power-law size-rank relationship* of the form $k\, d_k^{\alpha} = c$, where $c > 0$ and $\alpha > 0$ are constants, and where n_s determines the range of scaling [80]. Since this definition is simply a graph-specific version of (2.1) that allows for deviations from the power-law relationship for nodes with low connectivity, we again recognize that doubly logarithmic plots of d_k vs. k yield straight lines of slope $-\alpha$, at least for large d_k values.

This description of scaling degree sequence is general, in the sense that it applies to any given graph without regard to how it is generated and without reference to any underlying probability distributions or ensembles. That is, a scaling degree sequence is simply an ordered list of integers representing node connectivity and satisfying the scaling relationship of (2.1). In contrast, the current literature on complex networks focuses largely on *scaling degree distribution*, and thus a given degree sequence has the further interpretation as representing a realization of an iid sample of size n generated from a common scaling distribution of the type (2.2). This in turn is often induced by some random ensemble of graphs. This thesis will develop primarily a nonstochastic theory and thus focus on scaling degree sequences, but will clarify the role of stochastic models and distributions as well. In all cases, we will aim to be explicit about which is assumed to hold.

For graphs that are not trees, a first attempt at formally defining and relating the concepts of "scaling" or "scale-free" and "self-similar" through an appropriately defined notion of "scale invariance" is considered by Aiello Chung and Lu and described in [8]. In short, they view the evolution of a graph as a random process of growing the graph by adding new nodes and links over time. A model of a given graph evolution process is then called "scale-free" if "coarse-graining" in time yields scaled graphs that have the same power-law degree distribution as the original graph. Here "coarse-graining in time" refers to constructing scaled versions of the original graph by dividing time into intervals, combining all nodes born in the same interval into supernodes, and connecting the resulting

supernodes via a natural mapping of the links in the original graph. For a number of graph growing models, including the Barabási-Albert construction, Aiello Chung and Lu show that the evolution process is "scale-free" in the sense of being invariant with respect to time scaling (i.e., the frequency of sampling with respect to the growth rate of the model) and independent of the parameter of the underlying power-law node degree distribution (see [8] for details). Note that the scale invariance criterion considered in [8] concerns exclusively the degree distributions of the original graph and its coarse-grained or scaled counterparts. Specifically, the definition of "scale-free" considered by Aiello et al. is not "structural" in the sense that it depends on a macroscopic statistic that is largely uninformative as far as topological properties of the graph are concerned.

2.2 Scale-Free Networks

The development of graphic models for complex network topology can be traced back to 1959, when Erdös and Renyí [47] proposed random graph models and many properties were discovered in the limit of large graph size. In 1998 and 1999, small-world networks [120] and scale-free networks [20] were proposed separately and started an avalanche of work on modeling of complex networks by focusing on generating different variations of topology models to match the large-scale statistical properties. Scale-free networks, in particular, supposedly replicate empirically observed scaling node degree relationships that are not easily captured by traditional graph model like regular graphs or Erdös-Renyí random graphs [20], leading the trend of the pursuit of universal properties that transcend specific system details. It is in exactly what these properties are, and the theories to explain and exploit them, where big confusion arises.

We review the existing scale-free literature describing some of the most popular models and their most appealing features. This is then followed by a brief a critique of the existing theory of scale-free networks in general.

2.2.1 Main Scale-Free Properties

The main properties of scale-free graphs that appear in the existing literature can be summarized as

- scale-free networks have scaling (power-law) degree distribution.
- scale-free networks can be generated by certain random processes, the foremost among which is preferential attachment.
- scale-free networks have highly connected "hubs" which "hold the network together" and give the "robust yet fragile" feature of error tolerance but attack vulnerability.
- scale-free networks are generic in the sense of being preserved under random degree preserving rewiring.
- scale-free networks are self-similar.
- scale-free networks are universal in the sense of not depending on domain-specific details.

This variety of features suggest the potential for a rich and extensive theory. Unfortunately, it is unclear from the literature which properties are necessary and/or sufficient to imply the others, and if any implications are strict, or simply "likely" for an ensemble. Many authors apparently define scale-free in terms of just one property, typically scaling degree distribution or random generation, and appear to claim that some or all of the other properties are then consequences. In this part, we aim to clarify exactly what options there are in defining scale-free graphs and deriving their additional properties.

2.2.2 Existing Scale-Free Literature

We briefly review the existing treatment of the above properties, related historical results, and shortcomings of the current theory.

The ambiguity regarding the definition of "scale-free" originates with the original papers [20, 11], but have continued since. Here scale-free graphs appear to be defined both as

graphs with scaling or power-law degree distributions and as being generated by a stochastic construction mechanism based on *incremental growth* (i.e., nodes are added one at a time) and *preferential attachment* (i.e., nodes are more likely to attach to nodes that already have many connections). Indeed, the apparent equivalence of scaling degree distribution and preferential attachment, and the ability of thus-defined (if ambiguously so) scale-free network models to generate node degree statistics that are consistent with the ubiquity of empirically observed power-laws is the most commonly cited evidence that scale-free network mechanisms and structures are, in some sense, universal [11, 10, 19, 20, 22].

Models of preferential attachment giving rise to power-law statistics actually have a long history and are at least 80 years old. As presented by Mandelbrot [80], one early example of research in this area was the work of Yule [126], who in 1925 developed power-law models to explain the observed distribution of species within plant genera. Mandelbrot [80] also documents the work of Luria and Delbrück, who in 1943 developed a model and supporting mathematics for the explicit generation of scaling relationships in the number of mutants in old bacterial populations [77]. A more general and popular model of preferential attachment was developed by Simon [112] in 1955 to explain the observed presence of power-laws within a variety of fields, including economics (income distributions, city populations), linguistics (word frequencies), and biology (distribution of mutants in bacterial cultures). Substantial controversy and attention surrounded these models in the 1950s and 1960s [80]. A recent review of this history can also be found in [87]. By the 1990s, though, these models had been largely displaced in the popular science literature by models based on critical phenomena from statistical physics [16], only to resurface recently in the scientific literature in this context of "scale-free networks" [20]. Since then, numerous refinements and modifications to the original Barabási-Albert construction have been proposed and have resulted in scale-free network models that can reproduce power-law degree distributions with any $\alpha \in [1, 2]$, a feature that agrees empirically with many observed networks [10]. Moreover, the largely empirical and heuristic studies of these types of "scale-free" networks have recently been enhanced by a rigorous mathematical treatment that can be found in [26] and involves a precise version of the Barabási-Albert construction.

The introduction of scale-free network models, combined with the equally popular

(though less ambiguous) "small world" network models [120], reinvigorated the use of abstract random graph models and their properties (particularly node degree distributions) to study a diversity of complex network systems. For example, Dorogovtsev and Mendes [43] provide a "standard programme of empirical research of a complex network," which for the case of undirected graphs consist of finding (1) the degree distribution; (2) the clustering coefficient; (3) the average shortest-path length. The presumption is that these features adequately characterize complex networks. Through the collective efforts of many researchers, this approach has cataloged an impressive list of real application networks, including communication networks (the WWW and the Internet), social networks (author collaborations, movie actors), biological networks (neural networks, metabolic networks, protein networks, ecological and food webs), telephone call graphs, mail networks, power grids and electronic circuits, networks of software components, and energy landscape networks (again, comprehensive reviews of these many results are widely available [10, 19, 96, 43, 105]). While very different in detail, these systems share a common feature in that their degree distributions are all claimed to follow a power-law, possibly with different tail indices.

Regardless of the definitional ambiguities, the use of simple stochastic constructions that yield scaling degree distributions and other appealing graph properties represent for many researchers what is arguably an ideal application of statistical physics to explaining and understanding complexity. Since scale-free models have their roots in statistical physics, a key assumption is always that any particular network is simply a realization from a larger ensemble of graphs, with an explicit or implicit underlying stochastic model. Accordingly, this approach to understanding complex networks has focused on those networks that are most likely to occur under an assumed random graph model and has aimed at identifying or discovering macroscopic features that capture the "essence" of the structure underlying those networks. Thus preferential attachment offers a general and hence attractive "microscopic" mechanism by which a growth process yields an ensemble of graphs with the "macroscopic" property of power-law node degree distributions [21]. Second, the resulting scale-free topologies are "generic." Not only is any specific scale-free graph the generic or likely element from such an ensemble, but also *"... an important property of*

scale-free networks is that [degree preserving] random rewiring does not change the scale-free nature of the network" (see Methods Supplement to [59]). Finally, this ensemble-based approach has an appealing kind of "universality" in that it involves no model-specific domain knowledge or specialized "design" requirements and requires only minimal tuning of the underlying model parameters.

Perhaps most importantly, scale-free graphs are claimed to exhibit a host of startling "emergent" consequences of universal relevance, including intriguing self-similar and fractal properties (see below for details), small-world characteristics [15], and "hublike" cores. Perhaps the central claim for scale-free graphs is that they have highly connected hubs, what we term SF hubs, which "hold the network together." As noted, the structure of such networks is highly vulnerable (i.e., can be fragmented) to attacks that target these hubs [11]. At the same time, they are resilient to attacks that knock out nodes at random, since a randomly chosen node is unlikely to be a hub and thus its removal has minimal effect on network connectivity. In the context of the Internet, where scale-free graphs have been proposed as models of the router-level Internet [125], this has been touted "the Achilles' heel of the Internet" [11], a vulnerability that has presumably been overlooked by networking engineers. Furthermore, the hublike structure of scale-free graphs is such that the epidemic threshold is zero for contagion phenomena [104, 17, 106, 105], thus suggesting that the natural way to stop epidemics, either for computer viruses/worms or biological epidemics such as AIDS, is to protect these hubs [39, 28]. Proponents of this modeling framework have further suggested that the emergent properties of scale-free graphs contributes to truly universal behavior in complex networks [25] and that preferential attachment as well is a universal mechanism at work in the evolution of these networks [62, 43].

The scale-free story has successfully captured the interest and imagination of researchers across disciplines, and with good reason, as the proposed properties are rich and varied. Yet the existing ambiguity in its mathematical formulation and many of its most essential properties has created confusion about what it means for a network to be "scale-free." One possible and apparently popular interpretation is that scale-free means simply graphs with scaling degree *sequences*, and that this alone implies all other features listed above. We will show that this is incorrect, and in fact none of the features follows from scaling alone.

Even relaxing this to random graphs with scaling degree *distributions* is by itself inadequate to imply any further properties. One goal of this thesis is to clarify the reasons why these interpretations are incorrect, and propose minimal changes to fix them. The opposite extreme interpretation is that scale-free graphs are defined as having all of the above-listed properties. We will show that this is possible in the sense that the set of such graphs is not empty, but as a definition this leads to two further problems. Mathematically, one would prefer fewer axioms, and we will rectify this with a minimal definition. We will introduce a structural metric that provides a view of the extent to which a graph is scale-free and from which all the above properties follow, often with necessary and sufficient conditions. The other problem is that the canonical examples of apparent scale-free networks, the Internet and biological metabolism, are then very far from scale-free in that they have *none* of the above properties except perhaps for scaling degree distributions. This is simply an unavoidable conflict between these properties and the specifics of the applications, and cannot be fixed.

As a result, a rigorous theory of scale-free graphs must either define scale-free more narrowly than scaling degree sequences or distributions in order to have nontrivial emergent properties, and thus lose central claims of applicability, or instead define scale-free as merely scaling, but lose all the universal emergent features that have been claimed to hold for scale-free networks. We will pursue the former approach because we believe it is most representative of the spirit of previous studies and also because it is most inclusive of results in the existing literature. At the most basic level, simply to be a nontrivial and novel concept, scale-free clearly must mean more than a graph with scaling degree sequence or distribution. It must capture some aspect of the graph itself, and not merely a sequence of integers, stochastic or not, in which case the scale-free literature and this thesis would offer nothing new. Other authors may ultimate choose different definitions, but in any case, the results in this thesis clarify for the first time precisely what the graph theoretic alternatives are regarding the implications of any of the possible alternative definitions. Thus the definition of the word “scale-free” is much less important than the mathematical relationship between their various claimed properties, and the connections with real world networks.

2.3 Summary

In this chapter, we provide background knowledge on the scaling distribution and scale-free networks, which are essential for the exploration of complex networks.

To illustrate some key points about the existing claims regarding scale-free networks as adopted in the popular literature and their relationship with scaling degree distributions, we consider an application to the Internet where graphs are meant to model the Internet connectivity at the router-level in the next chapter.

Chapter 3

Internet Router-Level Topology

Previous studies on network topologies have focused on interpreting measurements or on phenomenological descriptions and evaluation of graph-theoretic properties of topology generators. Power-law degree distribution and scale-free networks dominate the current literature.

Arguing against this approach, and taking the Internet router-level topology as an example, we propose a complementary approach by emphasizing the network functionality together with practical constrains and trade-offs. We claim that very simple models that incorporate hard technological constraints on router bandwidth and connectivity, together with abstract models of user demand and network performance, can successfully depict and capture the intrinsic fundamentals of the Internet topology. In parallel, we provide evidence that scale-free networks as constructed by a constrained random process are inherently flawed.

This chapter is organized in the following manner.

We introduce the router-level topology and our work briefly in section 1 and we review the previous approaches to generating realistic Internet topologies in section 2. In section 3, we provide an alternate approach to understanding topology structure that explicitly incorporates router technology constraints, various economic constraints, and network performance at work in the construction of real networks. Then in section 4, we explore several real Internet topologies and show the consistence with our approach. In section 5, we discuss several performance related metrics for comparing and contrasting networks, particularly with the popular scale-free networks. We give examples and presents our findings

in section 6 and summarize our work at the end.

3.1 Introduction

A detailed understanding of the many facets of the Internet's topological structure is critical for evaluating the performance of networking protocols, for assessing the effectiveness of proposed techniques to protect the network from nefarious intrusions and attacks, or for developing improved designs for resource provisioning.

Recent attention on the large-scale topological structure of the Internet has been heavily focused on the *connectivity* of network components, whether they be machines in the router-level graph [55, 29] or entire subnetworks (Autonomous Systems) in the AS-level graph [52, 33]. A particular feature of network connectivity that has generated considerable discussion is the prevalence of heavy-tailed distributions in node *degree* (e.g., number of connections) and whether or not these heavy-tailed distributions conform to power-laws [48, 85, 34, 87]. This macroscopic statistic has greatly influenced the generation and evaluation of network topologies. In the current environment, degree distributions and other large-scale statistics are popular metrics for evaluating how representative a given topology is [56], and scale-free models and their different variations become the most popular models for the Internet topology [30, 63, 84, 125, 10, 9, 100].

Yet, from our viewpoint, this perspective is both incomplete and in need for corrective action. For one, there exist many different graphs having the *same distribution of node degree*, some of which may be considered *opposites* from the viewpoint of network engineering. Furthermore, there are a variety of distinctly different random graph models that might give rise to a given degree distribution, and some of these models may have no network-intrinsic meaning whatsoever. Finally, we advocate here an approach that is primarily concerned with developing a basic understanding of the observed high variability in topology-related measurements and reconciling them with the reality of engineering design. From this perspective, reproducing abstract mathematical constructs such as power-law distributions is largely a side issue.

In this chapter, we consider *a first-principles approach* to understanding Internet topol-

ogy at the *router-level*, where nodes represent routers and links indicate one-hop connectivity between routers. More specifically, when referring in the following to router-level connectivity, we always mean Layer 2, especially when the distinction between Layer 2 vs. Layer 3 issues is important for the purpose of illuminating the nature of the actual router-level connectivity (i.e., node degree) and its physical constraints. For router-level topology issues such as performance, reliability, and robustness to component loss, the physical connectivity between routers is more important than the virtual connectivity as defined by the higher layers of the protocol stack (e.g., IP, MPLS). Moreover, we use here the notion of "first-principles approach" to describe an attempt at identifying some *minimal* functional requirements and physical constraints needed to develop simple models of the Internet's router-level topology that are at the same time illustrative, representative, insightful, and consistent with engineering reality. Far from being exhaustive, this attempt is geared toward accounting for very basic network-specific aspects, but it can readily be enhanced if some new or less obvious functional requirements or physical constraints are found to play a critical role. Also, in the process of developing models of the Internet router-level connectivity that are "as simple as possible, but not simpler," we focus on single ISPs or ASes as the Internet's fundamental building blocks that are designed largely in isolation and then connected according to both engineering and business considerations.

While there are several important factors that contribute to the design of an ISP's router-level topology (e.g., available technology, economic viability, customer demands, redundancy and geography) and while opinions will vary about which and how many of these factors matter, we focus here on a few critical technological and economic considerations that we claim provide insight into the types of network topologies that are possible. In essence, we argue the importance of explicit consideration of the basic trade-offs that network designers must face when building real networks. In parallel, we provide evidence that scale-free network models of router-level connectivity whose construction is constrained by macroscopic statistics but is otherwise governed by randomness are inherently flawed. To this end, we introduce the notions of *network performance* as a new means for discerning important differences between generated and real network topologies. In so doing, we show that incorporating fundamental design details is crucial to the understand-

ing and evaluation of Internet topology.

3.2 Previous Work on Internet Topology

Two network topologies that have received significant attention from these experimental approaches are the *AS graph* (representing organizational interconnectivity between subnetworks) and the *router-level graph* of the Internet. Despite the challenges associated with the careful collection and interpretation of topology-related network measurements, significant efforts by the networking community are yielding an emerging picture of the large-scale statistical properties of these topologies [48, 55, 4, 29, 113, 115].

The development of abstract, yet informed, models for network topology evaluation and generation has followed the work of empiricists. The first popular topology generator to be used for networking simulation was the Waxman model [121], which is a variation of the classical Erdös-Rényi random graph [47]. The use of this type of random graph model was later abandoned in favor of models that explicitly introduce nonrandom structure, particularly hierarchy and locality, as part of the network design [40, 127]. The argument for this type of approach was based on the fact that an inspection of real networks shows that they are clearly not random but do exhibit certain obvious hierarchical features. This approach further argued that a topology generator should reflect the design principles in common use. For example, in order to achieve desired performance objectives, the network must have certain connectivity and redundancy requirements, properties which are not guaranteed in random network topologies. These principles were integrated into the Georgia Tech Internetwork Topology Models (GT-ITM).

These *structural topology generators* were the standard models in use until power-law relationships in the connectivity of both the AS-level and router-level graphs of the Internet were reported by Faloutsos et al. [48]. Since then, the identification and explanation of power-laws has become an increasingly dominant theme in the recent body of network topology literature [125, 34, 85, 124]. Since the GT-ITM topology generators fail to produce power-laws in node degree, they have often been abandoned in favor of new models

that explicitly replicate these observed statistics.[1] Examples of these generators include the INET AS-level topology generator [63], BRITE [84], BA [125], AB [10], GLP [30], PLRG [9], and the CMU power-law generator [100].

Each of the aforementioned degree-based topology generators uses one of the following three probabilistic generation methods. The first is *preferential attachment* [20] which says (1) the growth of the network is realized by the sequential addition of new nodes, and (2) each newly added node connects to some existing nodes preferentially, such that it is more likely to connect with a node that already has many connections. As a consequence, high-degree nodes are likely to get more and more connections resulting in a power-law in the distribution of node degree. For a precisely defined model that incorporates the key features of preferential attachment and is amenable to rigorous mathematical analysis, we refer to [26] and references therein. The second generation method is due to Chung and Lu [35] who considered a *general model of random graphs (GRG) with a given expected degree sequence*. The construction proceeds by first assigning each node its (expected) degree and then probabilistically inserting edges between the nodes according to a probability that is proportional to the product of the degree of the two given endpoints. If the assigned expected node degree sequence follows a power-law, the generated graph's node degree distribution will exhibit the same power-law. The third generation method, the *power-law Random Graph (PLRG)* [9], also attempts to replicate a given (power-law) degree sequence. This construction involves forming a set L of nodes containing as many distinct copies of a given node as the degree of that node, choosing a random matching of the the elements of L, and applying a mapping of a given matching into an appropriate (multi)graph.[2]

One of the most important features of networks that have power-law degree distributions and that are generated according to one of these probabilistic mechanisms is that they all tend to have a few centrally located and highly connected "hubs" through which essentially most traffic must flow. For the networks generated by preferential attachment, the central "hubs" tend to be nodes added early in the generation process. In the GRG model

[1] See however a comment by E. Zegura on router-level topology modeling, `http://www.caida.org/analysis/topology/ router-level-topology.xml`.

[2] It is believed that the PLRG and GRG models are "basically asymptotically equivalent, subject to bounding error estimates" [9].

as well as in the PLRG model, the nodes with high (expected) degree have higher probability to attach to other high degree nodes and these highly connected nodes form a central cluster. When using these models to represent the Internet, the presence of these highly connected central nodes in these networks has been touted its "Achilles' heel" because network connectivity is highly vulnerable to attacks that target the high-degree hub nodes [11]. It has been similarly argued that these high-degree hubs are a primary reason for the epidemic spread of computer worms and viruses [104, 28]. The presence of highly connected central nodes in a network having a power-law degree distribution is the essence of the *scale-free* network models, therefore we do not specifically differentiate degree-based models and scale-free models in this chapter.

However, this emphasis on power-laws and the resulting efforts to generate and explain them with the help of these degree-based methods have not gone without criticism. A widely known deficiency is that degree-based methods for topology generation produce merely descriptive models that are in general not able to provide correct physical explanations for the overall network structure [124]. The claim is that, in the absence of an understanding of the drivers of network deployment and growth, it is difficult to identify the causal forces affecting large-scale network properties and even more difficult to predict future trends in network evolution. Nevertheless, in the absence of concrete examples of such alternate models, degree-based methods have remained popular representations for large-scale Internet structure.

This chapter follows the previous arguments of [13] in favor of the need to explicitly consider the technical drivers of network deployment and growth. In spirit, it delivers for degree-based networks a similar message as [127] did for the random graph-type models [121] that were popular with networking researchers in the early 1990s. While [127] identified and commented on the inherent limitations of the various constructs involving Erdös-Rényi-type random graphs, our work points toward similar shortcomings and unrealistic features when working with probabilistic degree-based graphs.

3.3 A First-Principles Approach

A key challenge in using large-scale statistical features to characterize something as complex as the topology of an ISP or the Internet as a whole is that it is difficult to understand the extent to which any particular observed feature is "fundamental" to its structure. Here, we consider a complementary approach for thinking about network topology, in which we explore some of the practical constraints and trade-offs at work in the construction of real networks. In essence, we are asking the question, "What really matters when it comes to topology construction?" and argue that minimally one needs to consider the role of router technology and network economics in the network design process of a single ISP. The hope is that even a preliminary understanding of key factors, when combined with a more subtle use of statistics and graph theory, can provide a perspective that is more consistent both with observed measurements and the engineering principles at work in network design than with the current, at times conflicting, claims about the real Internet topology. In particular, given the current emphasis on the presence of power-laws in the connectivity of the router-level Internet, it is important to understand whether such variability is plausible, and if so, where it might be found within the overall topology. Fortunately, such an explanation is possible if one considers the importance of router technology and network economics in the design process.

3.3.1 Technology Constraints

In considering the physical topology of the Internet, a major constraint affecting the types of topologies available network designers is related to the routing equipment used to control the flow of traffic on the network, and the underlying *router technology constraints* are a significant force shaping network connectivity. Based on the technology used in the cross-connection fabric of the router itself, a router has a maximum number of packets that can be processed in any unit of time. This constrains the number of link connections (i.e., node *degree*) and connection speeds (i.e., bandwidth) at each router. This limitation creates a "feasible region" and corresponding "efficient frontier" of possible bandwidth-degree combinations for each router. That is, a router can have a few high bandwidth connections

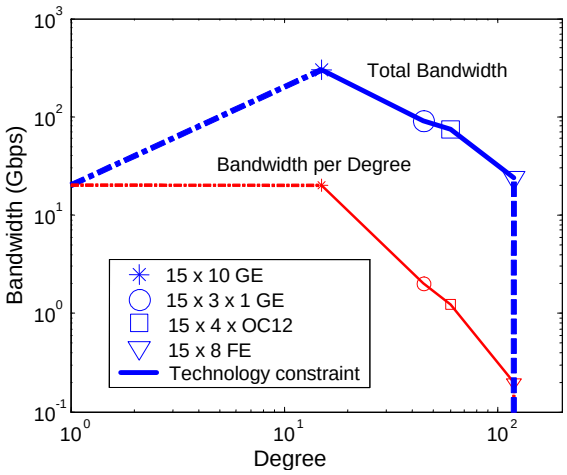

Figure 3.1: Technology constraint for Cisco 12416 Gigabit Switch Router(GSR): degree vs. bandwidth.

or many low bandwidth connections (or some combination in between). In essence, this means that routers must obey a form of *flow conservation* in the traffic that they can handle. While it is always possible to configure the router so that it falls below the efficient frontier (thereby underutilizing the router capacity), it is not possible to exceed this frontier (e.g., by having many high bandwidth connections).

Figure 3.1 shows the technology constraint for the Cisco 12416 GSR, which is one of the most expensive and highest bandwidth routers available from a 2002 Cisco product catalog [6]. Each point on the plot corresponds to a different combination of line cards and interfaces for the same router. This router has 15 available line card slots. When the router is configured to have less than 15 connections, throughput per degree is limited by the line-card maximum speed (10 Gbps) and the total bandwidth increases with the number of connections, while bandwidth per degree remains the same (dash-dot lines). When the number of connections is greater than 15, the total router bandwidth and bandwidth per degree decrease as the total number of connections increases (solid lines), up to a maximum of 120 possible connections for this router (dotted line). These three lines collectively define the feasible region for configuring this router.

Although engineers are constantly increasing the frontier with the development of new routing technologies, each particular router model will have a frontier representing its fea-

sible region, and network architects are faced with trade-offs between capacity and cost in selecting a router and then must also decide on the quantity and speed of connections in selecting a router configuration. Until new technology shifts the frontier, the only way to create throughput beyond the frontier is to build networks of routers.[3]

The current Internet is populated with many different router models, each using potentially different technologies and each having their own feasible region. However, these technologies are still constrained in their overall ability to trade-off total bandwidth and number of connections. Thus, networking products tend to be specialized to take advantage of one area of an aggregate feasible region, depending on their intended role within the network hierarchy. Consider an aggregate picture of many different technologies (shown in figure 3.2), used both in the network core and at the network edge. In addition to the Cisco 12000 GSR Series, the constraints on the somewhat older Cisco 7000 Series is also shown. Edge technologies are somewhat different in their underlying design, since their intention is to be able to support large numbers of end users at fixed (DSL, dial-up) or variable (cable) speeds. They can support a much greater number of connections (upwards of 10,000 for DSL or dial-up) but at significantly lower speeds. The shared access technology for broadband cable provides service comparable to DSL when the total number of users is about 100, but can only provide service equivalent to dial-up when the number of users is about 2000. Included also is the Linksys 4-port router, which is a popular LAN technology supporting up to 5100 MB Ethernet connections. Observe that the limits of this less expensive technology are well within the interior of the feasible region for core network routers. Collectively, these individual constraints form an overall aggregate constraint on available topology design.

We are not arguing that limits in technology fundamentally preclude the possibility of high-degree, high-bandwidth routers, but simply that the product offerings recently available to the marketplace have not supported such configurations. While we expect that companies will continue to innovate and extend the feasible region for router configuration, it remains to be seen whether or not the economics (including configuration and management)

[3]Recent product announcements from router manufacturers such as Juniper Networks, Avici Systems, and Cisco Systems suggest that the latest trend in technology development is to build scaleable multirack routers that do exactly this.

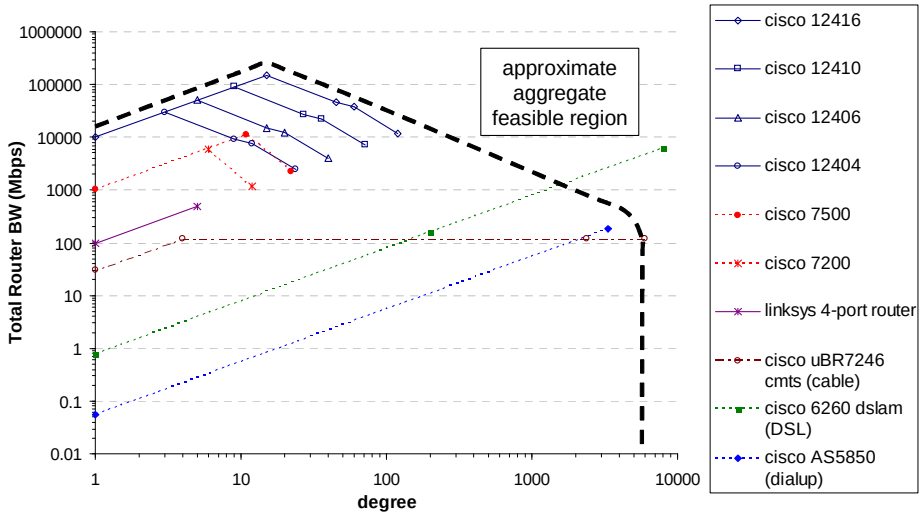

Figure 3.2: Aggregate picture of router technology constraints.

for these products will enable their wide deployment within the Internet.

3.3.2 Economic Considerations

Even more important than the technical considerations affecting router use are the economic considerations of network design and deployment, which are driven by customer demands and ultimately direct the types of technologies that are developed for use by network providers. For example, the cost of installing and operating physical links in a network can dominate the cost of the overall infrastructure, and since these costs tend to increase with link distance, there is tremendous practical incentive to design wired networks such that they can support traffic using the fewest number of links. The ability to share costs via multiplexing is a fundamental driver underlying the design of networking technologies, and the availability of these technologies enables a network topology in which traffic is aggregated at all levels of network hierarchy, from its periphery all the way to its core.

The development of these technologies has similarly followed the demands of customers, for whom there is wide variability in the willingness to pay for network bandwidths (Figure 3.3). For example, nearly half of all users of the Internet in North America still have dial-up connections (generally 56 kbps), only about 20% have broadband access (256 kbps-6 Mbps), and there is only a small number of users with large (10 Gbps) band-

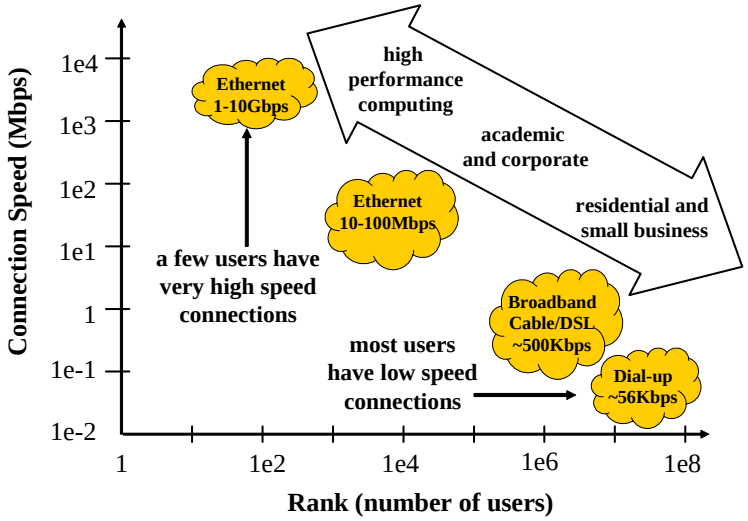

Figure 3.3: Aggregate picture of end user connection bandwidths for the Internet.

width requirements [12]. Again, the cost effective handling of such diverse end user traffic requires that aggregation take place as close to the edge as possible and is explicitly supported by a common feature that these edge technologies have, namely a special ability to support high connectivity in order to aggregate end user traffic before sending it towards the core.

The economic drive to minimize link costs promotes a topology that aggregates traffic as close to the network edge as possible. The use of multiplexing in a variety of routing technologies at the network edge supports this aggregation, and the wide variability in the bandwidth demands and geographies of end user connections suggests that one should expect wide variability in the measured connectivity of nodes at the network edge. Since it is generally accepted that most of the computers in the network are at its edge, it is reasonable to expect that the overall connectivity statistics of the network are dominated by those at the edge, which lead to the high variability in the whole Internet node connectivity.

3.3.3 Service Requirements

In addition to the constraints imposed by economic and router technology limitations, it is reasonable to expect that ISPs are driven to satisfy certain service requirements imposed by their customers or the industry at large. For example, most ISPs utilize service level

agreements (SLAs), which serve as business contracts with their major customers and their peers. SLAs typically specify terms such as delivered bandwidth and limits on service interruptions, and they often include financial penalties for failure to comply with their terms. While SLAs are often negotiated on an individual basis, competition among ISPs often creates industry norms that lead to standard SLA terms. Conversely, some ISPs use special terms in SLAs as a mechanism for differentiating their services and creating competitive advantages over rival companies.

From the provider's perspective, one simple metric for assessing whether or not a given network topology is "good" is its ability to handle the bandwidth requirements of its edge routers. We define *network performance* as the maximum throughput on the network under heavy traffic conditions based on a gravity model [128]. That is, we consider flows on all source-destination pairs of edge routers, such that the amount of flow X_{ij} between source i and destination j is proportional to the product of the traffic demand x_i, x_j at end points i, j, $X_{ij} = \alpha x_i x_j$, where α is some constant. We compute the maximum throughput on the network under the router degree bandwidth constraint,

$$\max_{\alpha} \quad \sum_{ij} \alpha x_i x_j \tag{3.1}$$

$$s.t. \quad RX \leq B, \tag{3.2}$$

where X is a vector obtained by stacking all the flows $X_{ij} = \alpha x_i x_j$ and R is the routing matrix (defined such that $R_{kl} = \{0, 1\}$ depending on whether or not flow l passes through router k). We use shortest path routing to get the routing matrix, and define B as the vector consisting of all router bandwidths according to the degree bandwidth constraint (Figure 3.2). Due to a lack of publicly available information on traffic demand for each end point, we assume the bandwidth demand at a router is proportional to the aggregated demand of any end hosts connected to it. In this manner, we allow for good bandwidth utilization of the higher level routers.[4] While other performance metrics may be worth considering, we claim that maximum throughput achieved using the gravity model provides a reasonable

[4]We also tried choosing the traffic demand according to other metric, such as the demand between routers is the same, or is as the product of their degrees as in [54], and qualitatively similar performance values are obtained but with different router utilization.

measure of the network to provide a *fair* allocation of bandwidth.

3.3.4 Heuristically Optimal Networks

Our objective is to develop a simple and minimal, yet plausible model for router-level topology that conforms to the technology constraints of routers, reflects link costs and high variability in end-user demand, and achieves reasonably "good" performance. We have argued that the perspective of an ISP in building a national scale network topology is driven by three factors. First, the need to minimize the long distance link costs means that it is driven to aggregate traffic from its edges to its core. Second, the design of its topology, particularly in the core, must conform to the technology constraints inherent in routers. Third, the network should have good performance, measured in terms of its ability to carry large volumes of traffic in a fair manner. While these are certainly not the only factors affecting design, we claim that these three drivers are a sensible starting point for understanding the relationship between ISP network design and resulting router-level topology.

Collectively, these constraints and functional requirements suggest that a "good" design is one in which individual links at the edge of the network have are aggregated in a manner such that the link capacities increase as one moves to the network core. In particular, edge routers may be connected to a large number of low bandwidth users or a smaller number of high bandwidth users. In contract, one can expect that the core is constructed as a loose mesh of high speed, low connectivity routers which carry heavily aggregated traffic over high bandwidth links. Accordingly, this meshlike core is supported by a hierarchical treelike structure at the edges whose purpose is to aggregate traffic through high connectivity. We will refer to this topology as *Heuristically Optimal Topology* (HOT) to reflect its consistency with real design considerations. By "heuristically" we have two meanings. On one side, we only consider the minimal set of functional requirements and constraints therefore the model reflects the most coarse grain level approximation. On the other side, our model does not exactly aim for the optimal solution, but any one that can provide a reasonable good performance. The optimal solution becomes less important since in the

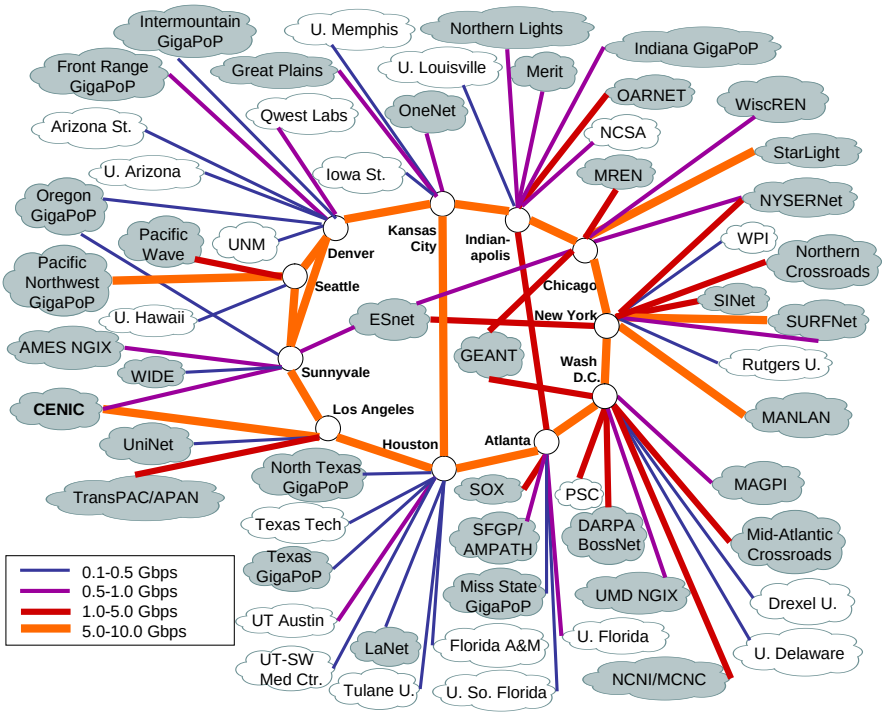

Figure 3.4: Abilene network. Each node represents a router, and each link represents a physical connection between Abilene and another network.

Internet, since much more trade-off between constraints and functions need to take into considerations and it is often impossible to find the best one.

3.4 Evidence from Real Internet

As evidence that this heuristic design shares similar qualitative features with the real Internet, we consider the real router-level connectivity of the Internet as it exists for the educational networks of Abilene and CENIC as well as a commercial tier-1 network from AT&T.

3.4.1 The Abilene Network

The Abilene Network (Figure 3.4) is the Internet backbone network for higher education, and it is part of the Internet2 initiative [1]. It is comprised of high-speed connections between core routers located in 11 U.S. cities and carries approximately 1% of all traffic in

North America.[5] The Abilene backbone is a sparsely connected mesh, with connectivity to regional and local customers provided by some minimal amount of redundancy. Abilene is built using Juniper T640 routers, which are configured to have anywhere from five connections (in Los Angeles) to twelve connections (in New York). Abilene maintains peering connections with other higher educational networks (both domestic and international) but does not connect directly to the commercial Internet.

3.4.2 The CENIC Network

Focusing in on a regional level, we consider California, where the Corporation for Education Network Initiatives in California (CENIC) acts as ISP for the state's colleges and universities [2]. Its backbone is similarly comprised of a sparse mesh of routers connected by high-speed links (Figure 3.5). Here, routing policies, redundant physical links, and the use of virtual private networks support robust delivery of traffic to edge campus networks. Similar observations are found when examining (where available) topology-related information of global, national, or regional commercial ISPs.

The CENIC backbone is comprised of two backbone networks in parallel—a high performance (HPR) network supporting the University of California system and other universities, and the digital California (DC) network supporting K-12 educational initiatives and local governments. Connectivity within each POP is provided by layer-2 technologies, and connectivity to the network edge is not shown. Each router has only a few high bandwidth connections, however each physical connection can support many virtual connections that give the appearance of greater connectivity to higher levels of the Internet protocol stack. ESnet and GEANT are other backbone networks.

In view of recent measurement studies [55, 4, 113], it is important to recognize that the use of technologies at layers other than IP will affect what traceroutelike experiments can measure. For example, the use of shared media at Layer 2 (e.g., Ethernet, FDDI rings) either at the network edge or at exchange points between ISPs can give the appearance of

[5] Of the approximate 80,000 - 140,000 terabytes per month of traffic in 2002 [99], Abilene carried approximately 11,000 terabytes of total traffic for the year [3]. Here, "carried" traffic refers to traffic that traversed an Abilene router. Since Abilene does not peer with commercial ISPs, packets that traverse an Abilene router are unlikely to have traversed any portion of the commercial Internet.

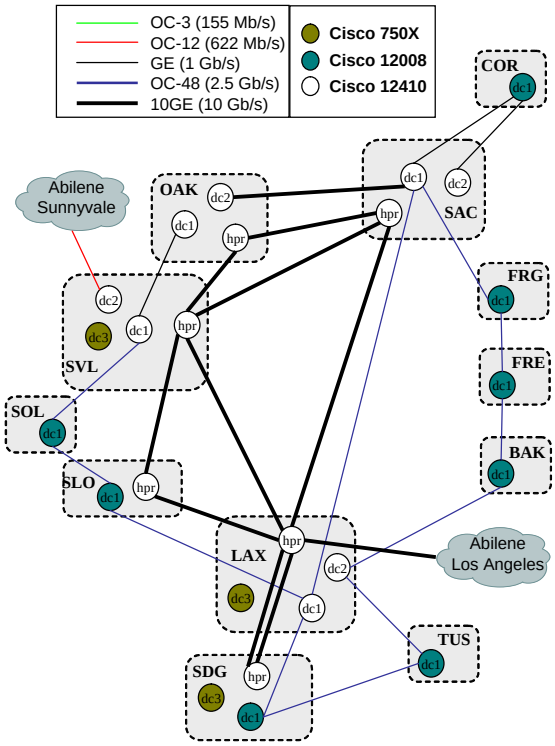

Figure 3.5: CENIC backbone.

high degree nodes. In an entirely different fashion, the use of Multiprotocol Label Switching (MPLS) at higher levels of the protocol stack can also give the illusion of one-hop connectivity at the lower layers when, in fact, there is none. Abilene is an ideal starting point for understanding heuristically optimal topologies, because within its backbone, there is no difference between the link layer topology and what is seen by IP. In contrast, the use of Ethernet and other link layer switching technologies within the CENIC POPs makes the interpretation and visualization of the physical intra-CENIC connectivity more difficult, but inferring the actual link layer connectivity is greatly facilitated by knowing the configurations of the individual CENIC routers as shown in figure 3.6. In the time since the Cisco catalog [6] was published, the introduction of a new line card (supporting 10x1GE interfaces) has shifted the feasible region for the model 12410 router. Since this router has nine available slots, this router can achieve a maximum of 90 Gbps with either nine 10 GE line cards or nine 10x1GE line cards. Although the shape of the feasible region may continue to change, its presence and corresponding implications for router configuration

and deployment will remain qualitatively the same.

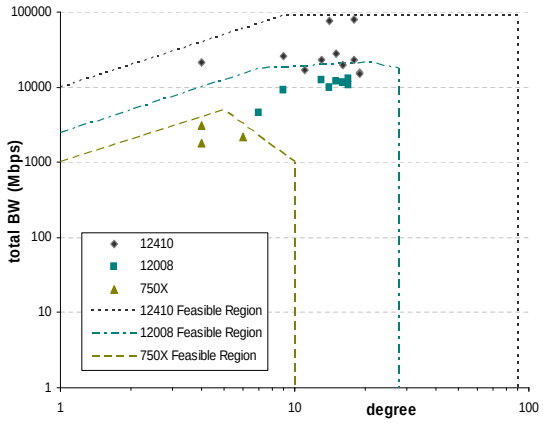

Figure 3.6: Configuration of CENIC routers.

3.4.3 A Tier-1 ISP Network

One of the overriding concerns of commercial ISPs in sharing topology data is that it will reveal information about its customers, thereby putting them at risk to competition. However, in cases where topology information is sufficiently anonymized and aggregated, we have found ISPs more willing to share and publish connectivity data. Here, we present aggregate router configuration information for AS 7018 (AT&T), as it existed during the second half of 2003. This Tier-1 ISP has hundreds of routers across the United States and is a major competitor in the national ISP market.

Figure 3.7 shows aggregate router configuration data for “core” and “access” routers in the ISP’s network as it existed during the second half of 2003. Routers are grouped into three different types: high-speed access routers, low-speed access routers, and core routers. For each group, we show the convex hull surrounding the points corresponding to the bandwidth-degree configuration for each router. Also shown is the feasible configuration region for a typical core router (i.e., the Cisco 12416 GSR) and a typical access router (i.e., the Cisco 7600). Here, “core routers” can be understood as those that provide long-haul connectivity between individual points of presence (PoPs) for the ISP. Conversely, “access routers” can be understood as those that provide aggregation connectivity between

the ISP and its customers within a PoP. For this ISP, access routers are further categorized according to whether they facilitate high-speed or low-speed connections. Although we are not able to show the configuration of individual routers for the ISP, we are able to present the convex hull containing the bandwidth-degree configuration for the routers of each type.

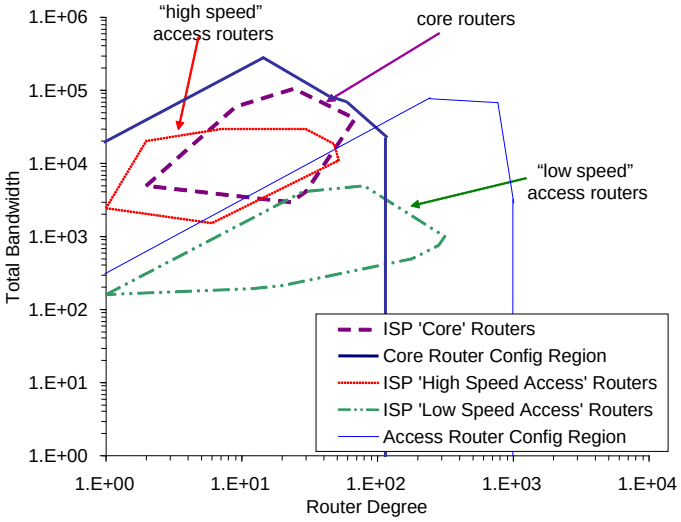

Figure 3.7: Configuration of a Tier-1 commercial ISP.

This aggregated information obscures individual router configurations as well as the total number of routers in each group, but it provides useful information nonetheless. First, the maximum number of connections to a core router is 68, while the maximum number of connections to a low-speed access router is 313. The maximum number of connections to a high-speed access router is less than that for both low-speed and core routers. Also, the relative position of these convex hulls reinforces the notion that routers are specialized according to their role (again, Figure 3.2). The core routers in this AS tend to have higher overall bandwidth than access routers, and they also tend to have fewer connections than many low-speed access routers. The high-speed access routers tend to have higher overall bandwidth but fewer connections than low-speed access routers. Also shown in figure 3.7 is the feasible region for representative core and access routers.[6] While certainly not all of the routers deployed in this AS were these specific router models, it is likely that some

[6]While the technology represented in the 2002 catalog [6] is now outdated, we argue that the product deployment lifecycle for routers makes it reasonable to believe that the production network of our ISP in 2003 is adequately reflected using this older technology.

of them were. Nonetheless, a striking feature of Figure 3.7 is the way in which the core routers conform rather tightly to the feasible configuration regions.

As a result of both educational networks topology and a commercial network configuration, we fully expect border routers to again have a few relatively high bandwidth physical connections supporting large amounts of aggregated traffic. In turn, high physical connectivity at the router level is expected to be firmly confined to the network edge.

3.5 Topology Metrics

3.5.1 Commonly-Used Metrics

Previous metrics to understanding and evaluating network topologies have been dominated by graph-theoretic quantities and their statistical properties, e.g., node-degree distribution, expansion, resilience, distortion and hierarchy [30, 56]. However we claim here that these metrics are inherently inadequate to capture the essential trade-offs of explicitly engineered networks.

Node degree distribution. In general, there are many networks having the same node degree distribution, as evidenced by the process of *degree-preserving rewiring*. This particular rewiring operation rearranges existing connections in such a way that the degrees of the nodes involved in the rearrangement do not change, leaving the resulting overall node degree distribution invariant. Accordingly, since the network can be rewired step-by-step so that the high degree nodes appear either at the network core or at its edges, it is clear that radically different topologies can have one and the same degree distribution (e.g., power-law degree distribution). In this fashion, degree-preserving rewiring is a means for moving within a general "space of network graphs," all having the same overall degree distribution.

Expansion, Resilience, Distortion. Introduced in [56], these metrics are intended to differentiate important aspects of topology. *Expansion* is intended to measure the ability of a node to "reach" other nodes within a given distance (measured by hops), *resilience* is intended to reflect the existence of alternate paths, and *distortion* is a graph theoretic metric that reflects the manner in which a spanning tree can be embedded into the topology. For

each of these three metrics, a topology is characterized as being either “Low” (L) or “High” (H). Yet, the quantitative values of expansion, resilience, and distortion as presented in [56] are not always easy to interpret when comparing qualitatively different topologies. For example, the measured values of expansion for the AS-level and router-level topologies show a relatively big difference (Figure 2(d) in [56]), however both of them are classified as “High,” suggesting that the degree-based generators compare favorably with measured topologies. In contrast, it could be argued that Tiers generates topologies whose expansion values match that of the measured router-level graph reasonably well (Figure 2(g) in [56]), but Tiers is classified to have “Low” expansion. Such problems when interpreting these metrics make it difficult to use them for evaluating differences in topologies in a consistent and coherent manner.

Nonetheless, these metrics have been used in [56] to compare measured topologies at the autonomous system (AS) level and the router level (RL) to topologies resulting from several generators, including degree-based methods (PLRG, BA, BRITE, BT, JinCJ00) and structural methods (GT-ITM’s Tiers and Transit-Stub), as well as several “canonical” topologies (e.g., random, mesh, tree, complete graph). It was observed that AS, RL, and degree-based networks were the only considered networks that share values “HHL” for expansion, resilience, and distortion respectively. Furthermore, of the canonical topologies, this “HHL” characterization was shared only by the complete graph (all nodes connected to each other). However, one canonical topology that was not considered was the “star” topology (i.e., having a single central hub), which according to their metrics would also be characterized as “HHL,” and which explains why the degree-based graphs (having high degree central hubs) fit this description. Yet, the fact that both a complete graph and a star could have the same characterization illustrates how this group of metrics is incomplete in evaluating network topology.

Hierarchy. For evaluating hierarchy, [56] considers the distribution of “link values,” which are intended to mimic the extent to which network traffic is aggregated on a few links (presumably, backbone links). However, the claim that degree-based generators, such as PLRG, do a better a job of matching the observed hierarchical features of measured topologies is again based on a qualitative assessment whereby previous structural genera-

tors (e.g., Tiers in GT-ITM) create hierarchy that is "strict" while degree-based generators result, like measured topologies, in hierarchies that are "moderate." This assessment is based on a model in which end-to-end traffic follows shortest path routes, however it also ignores any constraints on the ability of the network to simultaneously carry that end-to-end traffic.

These previous metrics appear to be inadequate for capturing what matters for real network topologies. Many of them lack a direct networking interpretation, and they all rely largely on qualitative criteria, making their application somewhat subjective. In what follows, we use the experience gained by these previous studies to develop metrics that are consistent with our first principles perspective. In particular, we consider several novel measures for comparing topologies that we show provide a minimal, yet striking comparison between degree-based probabilistic networks and networks inspired by engineering design.

3.5.2 Function-Related Metrics

Recognizing that the primary purpose for building a network is to carry effectively a projected overall traffic demand, we consider several means for evaluating the performance of the network.

Throughput. We define *network performance* as the maximum throughput on the network under heavy traffic conditions based on a gravity model [128]. This has been defined the same as the service requirement section 3.3.3.

$$Perf = \max_{\alpha} \sum_{ij} \alpha x_i x_j, \quad \text{such that } RX \leq B.$$

Router Utilization. In computing the maximum throughput of the network, we also obtain the total traffic flow through each router, which we term *router utilization*. Since routers are constrained by the feasible region for bandwidth and degree, the topology of the network and the set of maximum flows will uniquely locate each router within the feasible region. Routers located near the frontier are used more efficiently, and a router on the fron-

tier is saturated by the traffic passing through it. For real ISPs, the objective is clearly not to maximize throughput but to provide some service level guarantees (e.g., reliability), and modeling typical traffic patterns would require additional considerations (such as network overprovisioning) that are not addressed here. Our intent is not to reproduce real traffic, but to evaluate the raw carrying capacity of selected topologies under reasonable traffic patterns and technology constraints.

End User Bandwidth Distribution. In addition to the router utilization, each set of maximum flows also results in a set of bandwidths that are delivered to the end users of the network. While not a strict measure of performance, we consider as a secondary measure the ability of a network to support "realistic" end user demands.

Robustness to Failure. Another important issue in the design of ISP topologies is related to their reliability or robustness in the presence of equipment failure. Generally, network robustness is quantified in terms of the ability of the network to maintain its functionality after routers are removed and after rerouting of traffic. In the previous literature such as [11], this functionality is characterized by connectivity, which means size of largest connected cluster after removing failed nodes and related links from the network. Here we consider a simple metric more appropriate for Internet as the amount of original traffic (as measured by our previously defined notion of performance) that can still be served by the remaining network, possibly after some rerouting, but with the routers' bandwidth that remains its original value from feasible region.

3.6 Comparing Topologies

In this section, we compare and contrast the features of several different network graphs using the metrics described previously. Our purpose is to show that networks having the same (power-law) node degree distribution can (1) have vastly different features, and (2) appear deceivingly similar from a view that considers only graph theoretic properties.

3.6.1 A First Example

Our first comparison is made between five networks resulting from preferential attachment (PA), the GRG method with given expected node degree sequence, a generic heuristic optimal design, an Abilene-inspired heuristic design, and a heuristic suboptimal design. In all cases, the networks presented *have the same power-law degree distribution*. While some of the methods do not allow for direct construction of a selected degree distribution, we are able to use degree preserving rewiring as an effective (if somewhat artificial) method for obtaining the given topology. In particular, we generate the PA network first, then rearrange routers and links to get heuristically designed networks while keeping the same degree distribution. Lastly, we generate an additional topology according to the GRG method. What is more important here are the topologies and their different features, not the process or the particular algorithm that generated them.

Preferential Attachment (PA). The PA network is generated by following process: begin with 3 fully connected nodes, then in successive steps add one new node to the graph, such that this new node is connected to the existing nodes with probability proportional to the current node degree. Eventually we generate a network with 1000 nodes and 1000 links. Notice that this initial structure is essentially a tree. We augment this tree by successively adding additional links according to [10]. That is, in each step, we choose a node randomly and connect it to the other nodes with probability proportional to the current node degree. The resulting PA topology is shown in in figure 3.8(b) and has an approximate power-law degree distribution shown in figure 3.8(a).

General Random Graph (GRG) method. We use the degree sequence of the PA network as the expected degree to generate another topology using the GRG method. Notice that this topology generator is not guaranteed to yield a connected graph, so we pick the giant component of the resulting structure and ignore the self-loops as in [56]. To ensure the proper degree distribution, we then add degree one edge routers to this giant component. Since the total number of links in the giant component is generally greater than the number of links in an equivalent PA graph having the same number of nodes, the number of the edge routers we can add is smaller than in the original graph. The resulting topology is

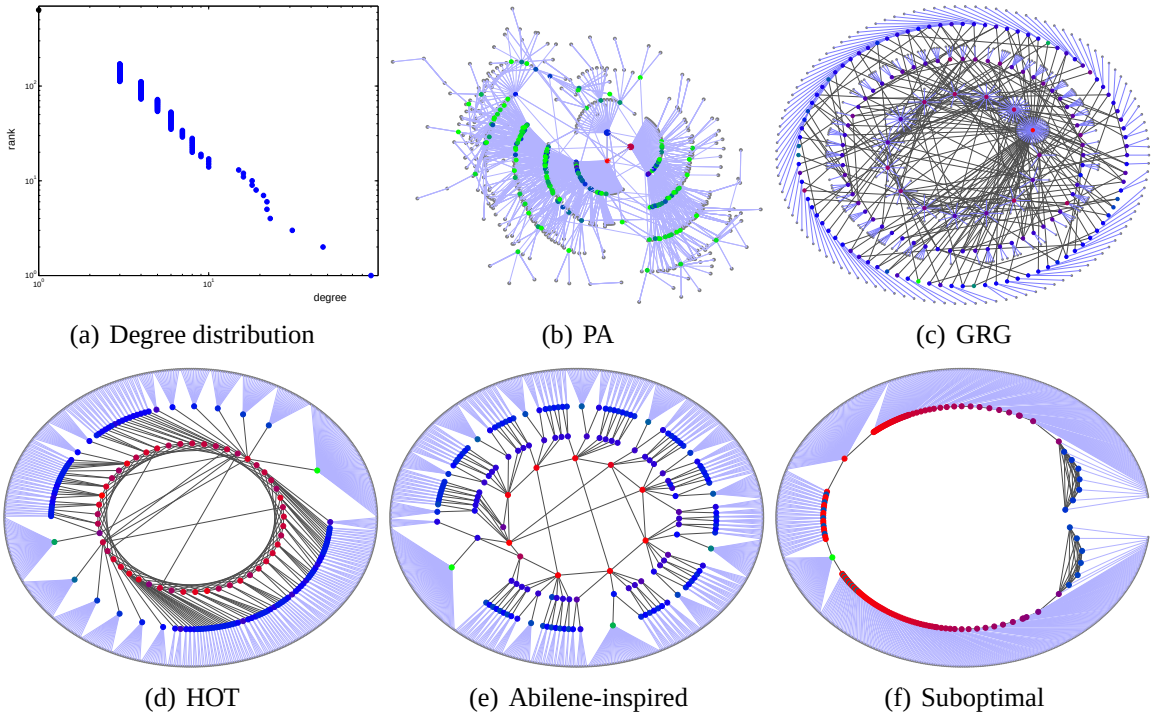

(a) Degree distribution (b) PA (c) GRG

(d) HOT (e) Abilene-inspired (f) Suboptimal

Figure 3.8: Five networks having the same node degree distribution as shown in sub-figure (a)

shown in figure 3.8(c), and while difficult to visualize all network details, a key feature to observe is the presence of highly connected central nodes.

Heuristically Optimal Topology (HOT). We obtain our HOT graph using a heuristic, nonrandom, degree-preserving rewiring of the links and routers in the PA graph. We choose 50 of the lower-degree nodes at the center to serve as core routers, and also choose the other higher-degree nodes hanging from each core as gateway routers. We adjust the connections among gateway routers such that their aggregate bandwidth to a core node is almost equally distributed. The number of edge routers placed at the edge of the network follows according to the degree of each gateway. The resulting topology is shown in figure 3.8(d). In this model, there are three levels of router hierarchy, each of which loosely correspond (starting at the center of the network and moving out toward the edges) to backbone, regional/local gateways, edge routers. Of course, several other “designs” are possible with different features. For example, we could have rearranged the network so as to have a different number of “core routers,” provided that we maintained our heuristic approach in using low-degree (and high bandwidth) routers in building the network core.

Abilene-inspired Topology. We claim that the backbone design of Abilene is heuristically optimal. To illustrate this, we construct a simplified version of Abilene in which we replace each of the edge network clouds in figure 3.4 with a single gateway router supporting a number of end hosts. We assign end hosts to gateway routers in a manner that yields the same approximate power-law in overall node degree distribution. The resulting topology with this node degree distribution is illustrated in figure 3.8(d).

Suboptimal Topology. For the purposes of comparison, we include a heuristically designed network that has not been optimized for performance (Figure 3.8(f)). This network has a chainlike core of routers, yet again has the same overall degree distribution.

Performance. For each of these networks, we impose the same router technological constraint on the nonedge routers. In particular, and to accommodate these simple networks, we use a fictitious router based on the Cisco GSR 12410, but modified so that the maximum number of ports it can handle coincides with the maximum degree generated above (see the dot-line in figure 3.9(b-f)). Thus, each of these networks has the same number of

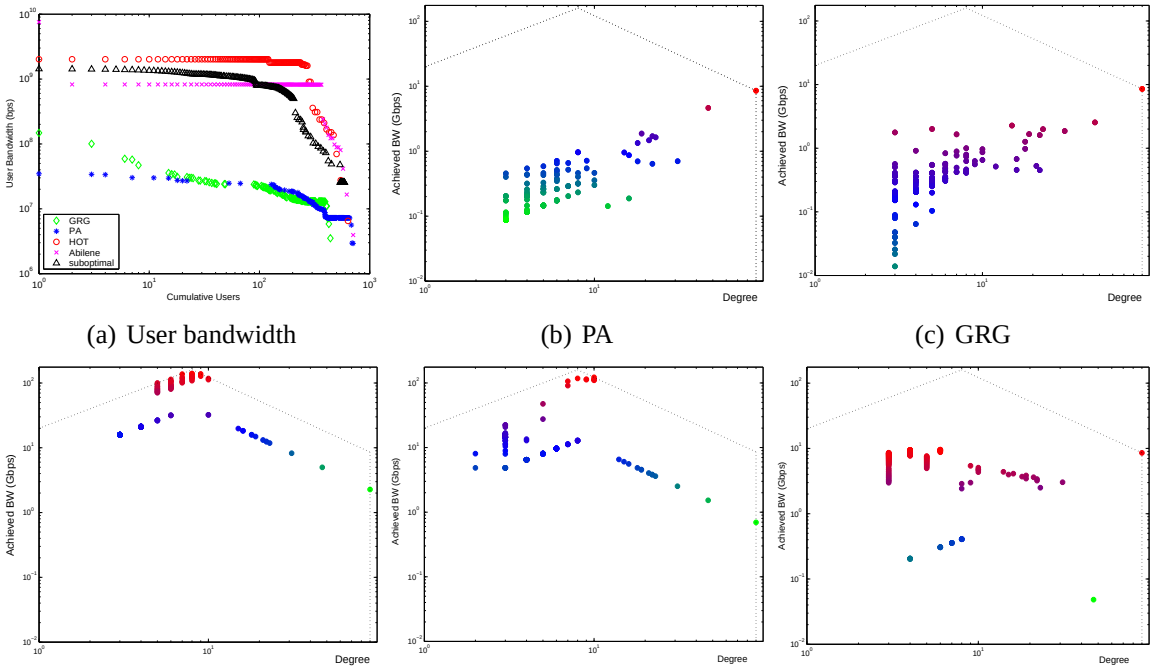

(a) User bandwidth (b) PA (c) GRG

(d) HOT (e) Abilene-inspired (f) Suboptimal

Figure 3.9: (a) Distribution of end user bandwidths for all the five different networks; (b)-(f) Router utilization for each network. The colorscale of a router on each plot differentiates its bandwidth which is consistent with the routers in figure 3.8.

nonedge nodes and links, as well as the same degree distribution among nonedge nodes. Collectively, these assumptions guarantee the same total "cost" (measured in routers) for each network. Using the performance index defined in section 3.5, we compute the performance of these five networks. Among the heuristically designed networks, the HOT model achieves 1130 Gbps and the Abilene-inspired network achieves 395 Gbps, while the suboptimal network achieves only 18.6 Gbps. For the randomly generated graphs, the PA and GRG achieve only 11.9 Gbps and 16.4 Gbps respectively, roughly 100 times worse than the HOT network. The main reason for PA and GRG models to have such terrible performance is exactly the presence of the highly connected "hubs" that create low-bandwidth bottlenecks. The HOT model's meshlike core, like the real Internet, aggregates traffic and disperses it across multiple high-bandwidth routers. We calculate the distribution of end user bandwidths and router utilization when each network achieves its best performance. Figure 3.9 (a) shows that the HOT network can support users with a wide range of bandwidth requirements, however the PA and GRG models cannot. Figure 3.9(d) shows that routers achieve high utilization in the HOT network, whereas, when the high degree "hubs" saturate in the PA and GRG networks, all the other routers are left under-utilized (Figure 3.9(b)(c)). The networks generated by these two degree-based probabilistic methods are essentially the same in terms of their performance.

Robustness to Failure. We use the PA model of Figure 3.8(b) and the HOT model of Figure 3.8(d) to compare network performance in the presence of successive router loss. Although robustness was not an explicit consideration in the heuristic construction of our HOT model, it has sufficient redundancy to merit this comparison, and the result is illustrative both of the extreme differences in these two types of models and of the manner in which our first-principles approach naturally allows for the incorporation of additional design considerations. Figure 3.10 shows the impact of deleting routers in succession from PA and HOT networks. We delete routers in succession from the PA and HOT networks, always targeting the worst-case router that has not yet been deleted. The measure of performance after deletion of a node is the amount of original traffic that can still be carried by the remaining network. Note that traffic can be rerouted, but the original router constraints remain intact.

Consistent with scale-free claims [11], the scale-free network is indeed fragile to the deletion of worst case nodes (here, worse case means highest degree); after removing the hubs, the performance drops by more than one order of magnitude. In contrast, the HOT networks is not only more robust to worst-case deletions (here, worst case are low-connectivity core nodes), but also shows high tolerance to deleting other nodes, particularly high-degree edge routers. In fact, because the scale-free network has such poor nominal performance to start with, it is worse intact than the HOT network after the latter has sustained substantial damage.

While a comprehensive study of large-scale network robustness is well beyond the scope of this thesis, our example illustrates two appealing features of the proposed first-principles approach. First, our detailed study of the technological and economic forces shaping the router-level topology of a single ISP provides convincing evidence that in today's Internet, the existence of highly connected routers in the core of the network is a myth. Size issues not notwithstanding, the real Internet is nothing like Figure 3.8 (b), but is qualitatively more like the network shown in figure 3.8(d): it cannot possibly have a hub-like core, and the highly connected nodes, if they exist, must be situated at the periphery of the network. Second, when trying to answer the question "What really matters when it comes to the ability of the Internet to perform in the presence of router or link losses?" we note that the ability of the network to "see damage and work around it" requires at a minimum adding some link redundancy (e.g., multihoming) and incorporating a simple abstraction of IP routing that accounts for the feedback mechanisms that react to the loss or failure of a network component. In particular, our approach makes it clear why the type of connectivity-only perspective pursued in [11] (i.e., one that completely ignores the existence of routing protocols sitting on top of the raw router-level connectivity) is bound to provide an overly simplistic and even misleading view of network robustness. Indeed, it is well known that the Internet's actual fragilities are not to physical attacks on routers or links, but to perturbations that were not part of the Internet's original design objectives [36], particularly misbehaving components or hijacked services.

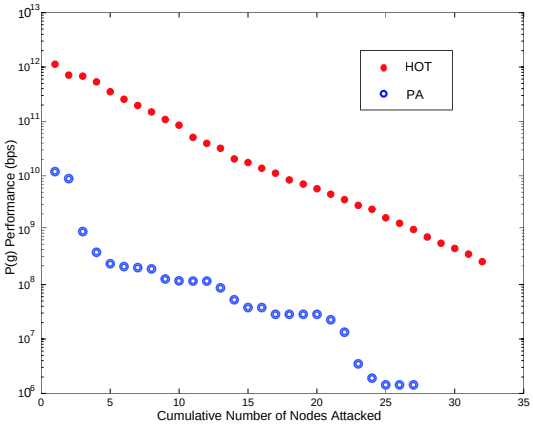

Figure 3.10: Robustness of network performance to router loss for the PA and HOT networks from Figure 3.8.

3.6.2 A Second Example

Figure 3.8 shows that graphs having the same node degree distribution can be very different in their structure, particularly when it comes to the engineering details. What is also true is that the same core network design can support many different end-user bandwidth distributions and that by and large, the variability in end-user bandwidth demands determines the variability of the node degrees in the resulting network. To illustrate, consider the simple example presented in figure 3.11, where the same network core supports different types of variability in end user bandwidths at the edge (and thus yields different overall node degree distributions). The network in figure 3.11(a) provides uniformly high bandwidth to end users; the network in figure 3.11(b) supports end user bandwidth demands that are highly variable; and the network in figure 3.11(c) provides uniformly low bandwidth to end users. Thus, from an engineering perspective, not only is there not necessarily any implied relationship between a network degree distribution and its core structure, there is also no implied relationship between a network's core structure and its overall degree distribution.

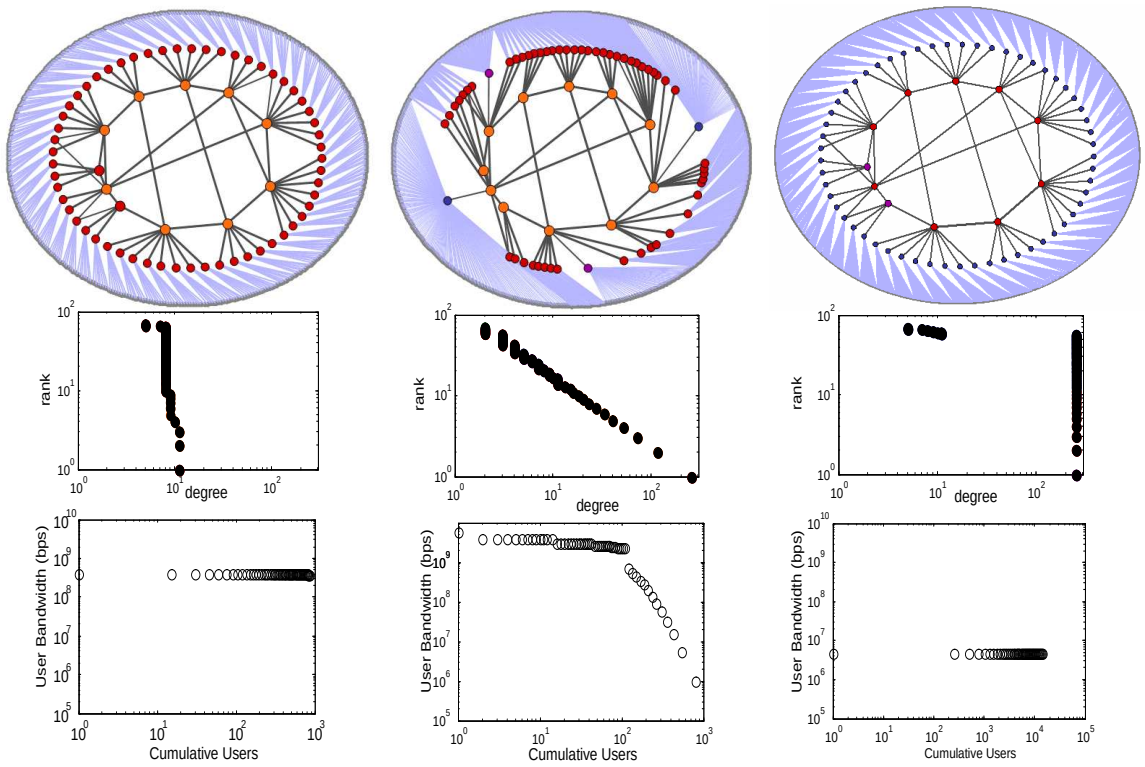

Figure 3.11: Distribution of node degree and end-user bandwidths for several topologies having the same core structure: (a) uniformly high bandwidth end users, (b) highly variable bandwidth end users, (c) uniformly low bandwidth end users.

3.7 Summary

The Internet router-level topology discussed in this chapter provides new insight into the space of all possible graphs that are of a certain size and are constrained by common macroscopic statistics, such as a given (power-law) node degree distribution. Scale-free models provide a relatively easy way to generate the desired power-law degree distribution, however their highly connected hubs have such bad performance as to make it completely unrealistic that they could reasonably represent a highly engineered system like an ISP network or the Internet as a whole. In contrast, we observe that even simple heuristically designed and optimized models that reconcile the trade-offs between link costs, router constraints, and user traffic demand result in configurations that have high performance and efficiency.

While the list of key factors of the router-level Internet considered by ISPs is far from

exhaustive in our model, what is striking is even simple domain-specific features shows how graphs that may be sensible from a connectivity-only perspective are no longer viable (e.g., nonrealizable or nonsensical) in the real world because of constraints that are imposed by their application domains.

This chapter focuses on the functionality of a complex network, particularly for the Internet router-level topology. We will introduce a structural metric for general complex networks, and elaborate its relationship to scale-free networks in the next chapter.

Chapter 4

The Structural Metric

In the previous chapter we have recognized that the power-law degree distribution and many aggregate statistics do not uniquely characterize a particular complex network. Many graphs with the same large-scale connectivity may have completely different structures and provide different performances. While the definition of functionality for different complex networks may vary, we introduce a structural metric, the *s*-metric, for general networks to characterize the highly connected hubs in their topologies, and show it is both necessary and useful for explaining the extreme differences among networks that have identical degree sequence, especially if it is scaling. This metric also yields considerable insight into the features of scale-free graphs. By focusing on a graph's structural properties and not on how it is generated, this approach does not depend on an underlying random graph model but is applicable to any graph of interest.

We organize this chapter as follows. In section 4.1, we investigate the definition and basic properties of the *s*-metric. We show how it relates to the "highly connected hubs" and the performance of in the Internet router-level topology. We also define the $s_{\max}$ and $s_{\min}$ graphs in both constrained and unconstrained graph space. We finally illustrate the fundamental relationship between the *s*-metric and diversity of degree sequences, the *s*-metric and the joint degree distribution. In section 4.2 and 4.3, we resolve the confusion of two important claims on scale-free networks: self-similarity and highly likely construction, and give their relationships to the *s*-metric. Section 4.4 provides insights into graph assortativity, a metric which directly relates to the *s*-metric, however is inappropriate for evaluating graphs due to the different background set in graph space. We conclude at the end.

4.1 Definitions and Properties

Let g be an undirected, simple, connected graph having $n = |\mathcal{V}|$ nodes and $l = |\mathcal{E}|$ links, where $\mathcal{V}$ and $\mathcal{E}$ are the sets of nodes and links, respectively. As before, define d_i to be the degree of node $i \in \mathcal{V}$, $D = \{d_1, d_2, \ldots, d_n\}$ to be the degree sequence for g, again assumed without loss of generality always to be ordered $d_1 \geq d_2 \geq \cdots \geq d_n$.

Within the space of all graphs having n nodes, let $\mathcal{G}(D)$ denote the considerably smaller subset of graphs having particular degree sequence D.

Not all sequences of integers D correspond to realizable graphs. One well-known characterization of whether or not a sequence D corresponds to a simple, connected graph is due to Erdös and Gallai [46], who observed that a sequence of positive integers $d_1, d_2, \ldots, d_n$ with $d_1 \geq d_2 \geq \cdots \geq d_n$ is *graphical* if and only if $\sum_{i=1}^{n} d_i$ is even and for each integer k, $1 \leq k \leq n-1$,

$$\sum_{j=1}^{k} d_j \leq k(k-1) + \sum_{j=k+1}^{n} \min(k, d_j).$$

The restriction to graphs having a particular degree sequence has been considered previously in the context of graph generation mechanisms [95, 35]. In particular, the Configuration Model (CM) [24, 89, 95] often serves as the null hypothesis of networks having a particular degree sequence, since it yields graphs that are maximally random (in the sense of maximum entropy) while conforming to a specified degree sequence D. In what follows, we will always restrict attention to graphs with a specified D.

In considering the structural features of a particular graph, we define, for any graph g having fixed degree sequence D, the s-metric

$$s(g) = \sum_{(i,j)\in\mathcal{E}} d_i d_j = \sum_{i\in\mathcal{V}} \sum_{j\in\mathcal{V}} \frac{1}{2}\, d_i a_{ij} d_j, \tag{4.1}$$

where $A = [a_{ij}]$ is the node adjacency matrix for the graph such that $a_{ij} = 1$ if nodes i, j are connected, $a_{ij} = 0$, otherwise. Accordingly, we assume without loss of generality that the number of nodes and links in the graph are represented by $n = |\mathcal{V}|$ and $l = |\mathcal{E}|$, respectively. Note that the summation in (4.1) is easily computed for any graph and does not depend on the process by which it was constructed.

Implicitly, the metric $s(g)$ measures the extent to which the graph g has a "hublike" core and is maximized when high-degree nodes are connected to other high-degree nodes. This observation follows from the *Rearrangement Inequality* [5], which states that if $a_1 \geq a_2 \geq \cdots \geq a_n$ and $b_1 \geq b_2 \geq \cdots \geq b_n$, then for any permutation $(a'_1, a'_2, \ldots, a'_n)$ of $(a_1, a_2, \ldots, a_n)$, we have

$$\begin{aligned} a_1b_1 + a_2b_2 + \cdots + a_nb_n &\geq a'_1b_1 + a'_2b_2 + \cdots + a'_nb_n & (4.2) \\ &\geq a_nb_1 + a_{n-1}b_2 + \cdots + a_1b_n. & (4.3) \end{aligned}$$

Since high $s(g)$-values are achieved only by connecting high-degree nodes to each other, and low $s(g)$-values are obtained by connecting high-degree nodes only to low-degree nodes, the s-metric moves beyond simple statements concerning the presence of "hub" nodes (as is true for any degree sequence D that has high variability) and attempts to quantify what role such hubs play in the overall structure of the graph. In particular, as we will show below, graphs with relatively high $s(g)$ values have a "hublike core" in the sense that these hubs play a central role in the overall connectivity of the network. We will also demonstrate that the metric $s(g)$ provides a view that is not only mathematically convenient and rigorous, but also practically useful as far as what it means for a graph to be "scale-free."

4.1.1 Performance vs. Structural Metrics

Before proceeding with a discussion of some of the features of the s-metric as well as for graphs having high $s(g)$ values, we revisit the toy models in figure 3.8 of Section 3.6.1 and consider the combined implications for the performance-oriented metric $Perf(g)$ and the structural metric $s(g)$. Figure 4.1 is a projection of g in simple and connected graph space onto a plane of $Perf(g)$ vs. $s(g)$ and will be useful throughout in visualizing the extreme diversity in the graph space. As structure always affects function, figure 4.1 shows a striking contrast which is observed by simultaneously plotting performance vs. the structural metric for these models. The HOT network has high performance and low s-metric while the PA and GRG networks have high s-metric but low performance. The interpretation of this

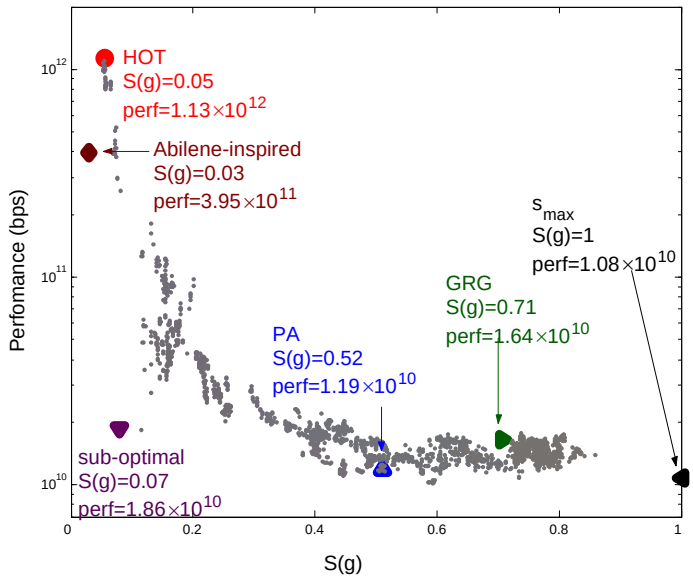

Figure 4.1: Performance vs. the structural metric for each topology, plus other networks having the same node degree distribution obtained by pairwise random rewiring of links. ($S(g) = (s(g) - s_{\min}) / (s_{\max} - s_{\min})$)

picture is that a careful design process explicitly incorporating technological constraints can yield high-performance topologies, but these are extremely far away from scale-free networks. In contrast, equivalent power-law degree distribution networks constructed by generic degree-based probabilistic constructions result in more highly connected hubs, but poor-performing topologies.

This viewpoint is augmented if one considers the process of pairwise random degree-preserving rewiring (details will be discussed in section 4.2.4) as a means to explore the space of graphs having the same overall degree distribution. In Figure 4.1, each point represents a different network obtained by random rewiring. Despite the fact that all of these graphs have the same overall degree distribution, we observe that a large number of these networks have relatively high *s*-metric and low performance. All of these graphs, including the PA and GRG networks, are consistent with "scale-free" models in the sense that they contain highly connected central hubs. The fact that there are very few high performance graphs in this space is an indication that it would be "hard" to find a relatively good design using random rewiring. We also notice that low *s*-metric itself does not guarantee a high performance network, as the network in figure 3.8(f) shows that it is possible to identify

small s and poorly performing networks. However, based on current evidence, it does appear to be the case that it is impossible using existing technology to construct a network that is both high performance and high s.

4.1.2 Joint Degree Distribution

As we see that the s-metric provides a good tool to differentiate graphs in the space of graphs with the same degree distribution, it is easy to justify this role from its relationship to the graph joint degree distribution (also called degree correlations). Given an appropriate statistical ensemble of graphs, the expectation of a random variable or random vector X is defined as

$$\langle X \rangle = \sum_{g \in G} X(g) P(g). \tag{4.4}$$

For example, for $1 \leq i \leq n$, let D_i be the random variable denoting the degree of node i for a graph $g \in G$ and let $D = \{D_1, D_2, \ldots, D_n\}$ be the random vector representing the node degrees of g. Then the *degree distribution* can be written in terms of an expectation of a random variable, namely

$$P(k) = \frac{1}{n} \left\langle \sum_{i=1}^{n} \delta[D_i - k] \right\rangle,$$

where

$$\delta[D_i(g) - k] = \begin{cases} 1 & \text{if node } i \text{ of graph } g \text{ has degree } k \\ 0, & \text{otherwise.} \end{cases}$$

We follow [43, Section 4.6] and define the joint degree distribution between two adjacent nodes having respective degree k and k' as follows.

Definition 4.1.1. *The joint degree distribution between two neighbors having degrees k and k' is defined by*

$$P(k, k') = \frac{1}{n^2} \left\langle \sum_{i,j=1}^{n} \delta[d_i - k] a_{ij} \delta[d_j - k'] \right\rangle, \tag{4.5}$$

where the a_{ij} are elements of the network node adjacency matrix such that

$$a_{ij} = \begin{cases} 1 & \textit{if nodes } i, j \textit{ are connected}, \\ 0 & \textit{otherwise}, \end{cases}$$

and where the random variables $\delta[D_i - k]$ *are as above.*

As an expectation of indicator-type random variables, $P(k, k')$ can be interpreted as the probability that a randomly chosen link connects nodes of degrees k and k', therefore $P(k, k')$ is also called the "degree-degree distribution" for links. Observe that for a given graph g having degree sequence D,

$$\begin{aligned} s(g) &= \sum_{(i,j)\in\mathcal{E}} d_i d_j \\ &= \sum_{(i,j)\in\mathcal{E}} \sum_{k\in D} k\delta[d_i - k] \sum_{k'\in D} \delta[d_j - k']k' \\ &= \sum_{(i,j)\in\mathcal{E}} \sum_{k\in D} \sum_{k'\in D} k\delta[d_i - k]\delta[d_j - k']k' \\ &= \frac{1}{2} \sum_{k,k'\in D} kk' \sum_{i,j=1}^{n} \delta[d_i - k]a_{ij}\delta[d_j - k']. \end{aligned}$$

Thus, there is an inherent relationship between the structural metric $s(g)$ and the joint degree distribution, which we formalize as follows.

Proposition 4.1.1.

$$\langle s\rangle = \frac{n^2}{2} \sum_{k,k'} kk' P(k, k'). \tag{4.6}$$

Proof. For fixed degree sequence D,

$$\begin{aligned} \langle s\rangle &= \left\langle \frac{1}{2} \sum_{k,k'\in D} kk' \sum_{i,j=1}^{n} \delta[d_i - k]a_{ij}\delta[d_j - k'] \right\rangle \\ &= \frac{1}{2} \sum_{k,k'\in D} kk' \left\langle \sum_{i,j=1}^{n} \delta[d_i - k]a_{ij}\delta[d_j - k'] \right\rangle \\ &= \frac{n^2}{2} \sum_{k,k'\in D} kk' P(k, k'). \end{aligned}$$

□

This result shows that for an ensemble of graphs having degree distribution D, the expectation of s can be written purely in terms of the aggregation of joint degree distribution.

As graphs being constrained by the same degree distribution, which is the first order statistics of the graph space, joint degree distribution provides the flexibility on the second order statistics in the space. Therefore, it is not hard to see that the *s*-metric, as the aggregation of the joint degree distribution, plays the most important role to differentiate graphs with the same degree distribution.

It immediately follows that even for the graphs with the same joint degree distribution, the third order degree distribution varies and the *s*-metric cannot tell the difference. Certainly, our methodology could enable us to expand the *s*-metric in terms of the aggregation of third order degree correlation. Here we just use *s*-metric to argue attention should not only devote to degree sequences in order to measure the structure of complex networks. It is clear that such sequences alone are insufficient to characterize the aggregate structure of a graph, and the graph space is extremely diverse. One can imagine that the space of graph is shrinking when we put constraint on higher order of degree distribution. However, the higher order of degree distribution, the more difficult it is to characterize [78].

4.1.3 The s_{max} and s_{min} Graphs

In general, the set $\mathcal{G}(D)$ will have many elements exhibiting a range of *s*-values. Within this space, we define the s_{max} and s_{min} graphs within $\mathcal{G}(D)$ as those having the maximum and minimum *s* values, respectively. To facilitate the derivation of these values, we introduce the vector

$$Z \equiv \{\overbrace{d_1,\ldots,d_1}^{d_1 \text{ elements}},\overbrace{d_2,\ldots,d_2}^{d_2 \text{ elements}},\ldots,\overbrace{d_n,\ldots,d_n}^{d_n \text{ elements}}\}, \tag{4.7}$$
$$\underbrace{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}_{\sum_{i=1}^{n} d_i \text{ elements}}$$

which is simply derived from the original degree sequence *D*. The $s_{\max}$ and $s_{\min}$ values within $\mathcal{G}(D)$ can be described in terms of *Z* in the following manner. Since $\mathcal{G}(D)$ only requires its elements to satisfy the degree sequence *D* (and ignores issues such as connectedness, multiple links, etc.). According to equation (4.2), it is easy to show that within $\mathcal{G}(D)$, one has

$$s_{\max} \leq \frac{1}{2} ZZ^T, \tag{4.8}$$

with equality achieved in practice only under certain circumstances (e.g., when the elements of D are all even or there is an even number of elements having any particular odd value). Accordingly, it follows that

$$s_{\min} \geq \frac{1}{2} Z\hat{Z}^T, \tag{4.9}$$

where $\hat{Z}$ is simply the vector Z with elements in reverse order. However, unlike the case in (4.8) where equality is achieved in practice only sometimes and the actual value may deviate considerably from the upper bound, the relationship in (4.9) holds with approximate equality and typically the $s_{\min}$ value deviates from the lower bound by only a single pair of links, if at all.

It is easy to see that the $s_{\max}$ value can be rewritten as

$$s_{\max} \approx \sum_{i=1}^{n} (d_i/2) \cdot (d_i)^2 = \sum_{i=1}^{n} (d_i)^3/2, \tag{4.10}$$

which is achieved in effect by creating primarily self-loops among the nodes in the network and then connecting remaining "stubs" in order of decreasing d_i (see appendix A.1 for details). To the best of our knowledge, there does not exist a comparable analytic formula (or interpretation) for the $s_{\min}$ graph in $\mathcal{G}(D)$.

Many graphs of practical interest have additional conditions imposed by functional or domain constraints, such as a requirement to be connected or a restriction against self-loops or multiple connections. Thus, in our investigation we also consider the restricted set of all *simple and connected graphs* having the same degree sequence D, which we denote as $G(D)$. Note that $G(D) \subset \mathcal{G}(D)$ and that most randomly generated graphs with particular D will be neither simple nor connected, so this is an important and nontrivial restriction. From these definitions it follows that

$$\frac{1}{2} Z\hat{Z}^T \leq s_{\min}^{\mathcal{G}(D)} \leq s_{\min}^{G(D)} \leq s_{\max}^{G(D)} \leq s_{\max}^{\mathcal{G}(D)} \leq \frac{1}{2} ZZ^T.$$

Although bounding values for the min and max elements of $\mathcal{G}(D)$ can be directly obtained from equations (4.8) and (4.9), obtaining $s_{\max}$ and $s_{\min}$ values within the restricted space

$G(D)$ is more complicated.

Given a particular degree sequence D, it is possible to use a deterministic procedure in order to construct the $s_{\max}$ graph in $G(D)$. The details of this construction procedure are presented in appendix A.2, but the basic idea is to order all potential links (i, j) for all $i, j \in \mathcal{V}$ according to their *weight* $d_i d_j$ and then add them one at a time in a manner that results in a simple, connected graph having degree sequence D. While simple enough in concept, this type of "greedy" heuristic procedure may have difficulty achieving the intended sequence D due to the global constraints imposed by connectivity requirements, but it works well in practice for most graphs (again, see A for details). Obtaining the $s_{\min}$ value is less exact, and it is easy to show that the $s_{\min}$ graph is not unique. Whitney and Alderson [122] have recently used a heuristic approach, originally proposed by Maslov and Sneppen [81], which employs a Metropolislike algorithm based on successive rewiring to obtain $s_{\min}$ values within $G(D)$. Unfortunately, this method is inefficient and does not reliably obtain the actual $s_{\min}^{G(D)}$ value. However, in practice one finds that $\frac{1}{2} Z\hat{Z}^T \approx s_{\min}^{\mathcal{G}(D)} \approx s_{\min}^{G(D)}$, so in the remainder of this chapter we use the $s_{\min}^{\mathcal{G}(D)}$ value defined in (4.9), as an approximate (and more conservative) bounding value for $s_{\min}^{G(D)}$.

4.1.4 Diversity of Degree Sequence

As a measure of graph structure, the s-metric provides a simple means for contrasting the differences between graphs having the same degree sequence, and in this chapter we use it exclusively as a means for measuring the *diversity* within this particular space of graphs. In particular, the extreme points $s_{\max}$ and $s_{\min}$ serve as meaningful reference points for individual graphs and the space as a whole, and for a given D the difference $s_{\max} - s_{\min}$ provides a measure of how different the absolute extremes are. Using this perspective, it is not hard to see that the amount of diversity for graphs having a particular D is related to the amount of *variability* within the sequence D itself. Following Section 2.1.1, we characterize variability with the standard measure of *(sample) coefficient of variation (CV)*,

which for a given sequence $D = (d_1, d_2, \ldots, d_n)$ is defined as

$$CV(D) = \sigma(D)/\langle d \rangle, \tag{4.11}$$

where $\langle d \rangle = n^{-1} \sum_{k=1}^{n} d_k$ is the average node degree, and we measure deviations of the d_i from its average $\langle d \rangle$ using the sample standard deviation, $\sigma(D) = (\sum_{k=1}^{n} (d_k - \langle d \rangle)^2 / (n-1))^{1/2}$.

For graphs with regular structure that have low variability in their degree sequence D, there is typically very little diversity in the corresponding space of graphs $G(D)$. Consider as an extreme example, a 1 dimensional lattice (i.e., a *chain*) with the degree sequence $D_{chain} = \{2, 2, 2, \ldots, 2, 1, 1\}$. One can easily show that for a chain consisting of n nodes,

$$CV(D_{chain}) = \frac{n^{1/2}(n-2)^{1/2}}{2^{1/2}(n-1)^{3/2}}$$

and thus $CV(D_{chain}) \to 0$ as $n \to \infty$. It is easy to see that there is no diversity among graphs having degree sequence D_{chain}, since all n-node chains are isomorphic to one another in $G(D)$ and thus $s_{\min} = s_{\max}$.

For sequences D with increasing $CV(D)$, graph diversity as measured by the range $s_{\max} - s_{\min}$ also increases. Here, we leverage two classes of graphs as reference points. For graphs with a degree sequence having an exponential form, $k\, e^{\lambda d_k} \approx c$ for constant $c > 0$ (denoted here as D_{exp}), one observes that $CV(D_{exp}) \to \kappa$ (a constant) as $n \to \infty$. In contrast, the *scale-free* graphs [20] exhibit divergent CV. It is easy to show that degree sequences $D_{scaling}$ with $\alpha < 2$ follow $CV(D_{scaling}) \to \infty$ as $n \to \infty$. As we will show below, these classes of graphs yield degree sequences with measurably different levels of diversity.

Although one might expect that graph diversity simply increases with $CV(D)$, this need not be the case. Consider a *star* consisting of a single central node that connects to all others and having degree sequence $D_{star} = \{n-1, 1, 1, \ldots, 1\}$. One can similarly show that

$$CV(D_{star}) = \frac{n^{1/2}(n-2)}{2(n-1)},$$

and thus $CV(D_{star}) \to \infty$ as $n \to \infty$. However, like the chain, there is no diversity among graphs having degree sequence D_{star} (i.e., all stars are isomorphic to one another in $G(D)$

and $s_{\min} = s_{\max}$).

In order to make the previous discussion more concrete, we now consider a simple experiment to investigate the role of $CV(D)$ in determining the diversity for graphs having particular D. For purposes of exposition, we begin with a study of acyclic graphs (i.e., trees) and then later comment on how our results apply to general graphs. Our experiment uses incremental growth via preferential attachment as described in [21], in which each newly added node connects to an existing node k with probability

$$\Pi(k) = b\frac{(d_k)^\gamma}{\sum_j (d_j)^\gamma}, \tag{4.12}$$

where d_k is again the degree of node k, and γ is a parameter that tunes the attachment mechanism. The resulting graph is simple and connected, thus an element of $G(D)$, although the degree sequence D that is realized will vary from trial to trial. Clearly, $\gamma = 0$ is equivalent to uniform attachment (resulting in D_{exp}), while $\gamma = 1$ is equivalent to linear preferential attachment used in the Barabási-Albert model [10] (resulting in $D_{scaling}$). A similar type of model was also considered in [67]. Note also that as $\gamma \to \infty$ each newly added node attaches to the maximum degree node (resulting essentially in D_{star}), while as $\gamma \to -\infty$ each newly added node attaches to the minimum degree node (resulting essentially in D_{chain}). In what follows, we first restrict attention to the case where $b = 1$ (i.e., we generate acyclic graphs) and consider a range of values for γ in order to generate graphs having a variety of degree sequences.

Figure 4.2 shows the result of an experiment in which for each trial we generate a tree having $n = 100$ nodes using preferential attachment rule given by equation (4.12). That is, each trial results in a tree having its own degree sequence D and s-value. In generating these graphs, we use various attachment exponents γ, but only for the purpose of realizing graphs with a diversity of degree sequences. In what follows we focus primarily on the degree sequence D and the constraints it places on the space of graphs, not the attachment exponent γ that led to D. For each degree sequence D, we then calculate $CV(D)$ as well as the corresponding $s_{\max}$ and $s_{\min}$ values as described above. The resulting picture in figure 4.2(a) shows a striking relationship between $CV(D)$ and the range of possible s-

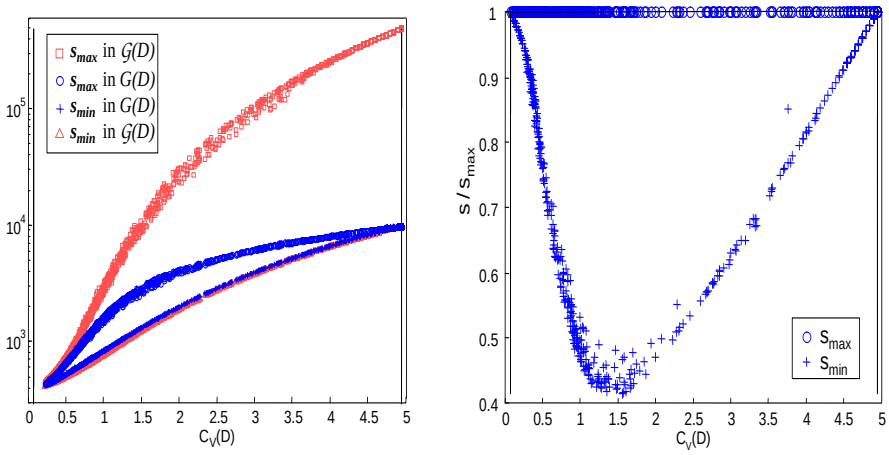

(a) (b)

Figure 4.2: Graph diversity and degree sequence diversity. (a) The $s_{\min}$ and $s_{\max}$ values in both $\mathcal{G}(D)$ and $G(D)$ vs. the $CV(D)$ of the corresponding degree sequence. Note that $s_{\min}^{G(D)} \approx s_{\min}^{\mathcal{G}(D)}$. (b) The $s_{\min}$ and $s_{\max}$ in $G(D)$, each normalized by their respective $s_{\max}$.

values. One observes that while the $s_{\max}$ and $s_{\min}$ values increase with $CV(D)$ for both the unconstrained space $\mathcal{G}(D)$ and the constrained space $G(D)$, the differences given by $s_{\max} - s_{\min}$ for each space behave differently at the maximal values of $CV(D)$. Specifically, this difference within the unconstrained space $\mathcal{G}(D)$ increases with $CV(D)$, but it is zero at *both* extremes of $CV(D)$ for the simple, connected graphs in $G(D)$ (again, the limiting cases of a chain and a star). It is also worth noting that the values for $s_{\min}^{G(D)}$ and $s_{\min}^{\mathcal{G}(D)}$ are so close as to be indistinguishable, further supporting our choice to treat these values as equivalent. Figure 4.2(b) presents the same information for $s_{\max}$ and $s_{\min}$ within $G(D)$, but normalizes the s-values for each graph against its respective $s_{\max}$ value, thus resulting in a feasible range $[0, 1]$ for each graph. Collectively, this suggests that for a given degree sequence one needs "enough" variability to enable diversity among simple, connected graphs but that "too much" variability actually becomes a constraint within the space $G(D)$, something that Maslov et al. [82] have described as essentially a finite size effect.

Although it is now well understood that there can be many graphs having the same degree sequence and that these graphs may have considerable structural differences, quantifying these differences and their implications in terms of real systems remains the topic of active research. Taken by itself, this observation is neither groundbreaking nor surprising. For some time, there has been a general recognition in the literature that the degree

sequence of a graph can provide only a simplistic characterization of its properties, and this has led many researchers to consider more sophisticated descriptions of graph structure.

4.2 The s-metric and Self-Similarity

When viewing graphs as multiscale objects, natural transformations that yield simplified graphs are pruning of nodes at the graph periphery or collapsing of nodes, although these are only the simplest of many possible "coarse-graining" operations that can be performed on graphs. These transformations are of particular interest because they are often inherent in measurement processes that are aimed at detecting the connectivity structure of actual networks. We will use these transformations to motivate that there is a plausible relationship between high-$s(g)$ graphs and self-similarity, as defined by these simple operations. We then consider the transformation of random pairwise degree-preserving (link) rewiring that suggests a more formal definition of the notion of a self-similar graph.

4.2.1 Graph Trimming by Link Removal

Here, we consider the properties of $s_{\max}$ graphs under the operation of graph trimming, in which links are removed from the graph one at a time. Recall that by construction, the links in the $s_{\max}$ graph are selected from a list of potential links (denoted as (i, j) for $i, j \in \mathcal{V}$) that are ordered according to their weights $d_i d_j$. Denote the (ordered) list of links in the $s_{\max}$ graph as $\mathcal{E} = \{(i_1, j_1), (i_2, j_2), \ldots, (i_l, j_l)\}$, and consider a procedure that removes links in reverse order, starting with (i_l, j_l). Define $\tilde{g}_k$ to be the remaining graph after the removal of all but the first $k-1$ links, (i.e., after removing $(i_l, j_l), (i_{l-1}, j_{l-1}), \ldots, (i_{k+1}, l_{k+1}), (i_k, l_k)$). The remaining graph will have a partial degree sequence $\tilde{D}_k = \{d_1', d_2', \ldots, d_k'\}$, where $d_m' \leq d_m, m = 1, 2, \ldots, k$, but the original ordering is preserved, i.e., $d_1' \geq d_2' \geq \cdots \geq d_k'$. This last statement holds because when removing links starting with the smallest $d_i d_j$, nodes will "lose" links in reverse order according to their node degree.

Observe for trees that removing a link is equivalent to removing a node (or subtree), so we could have equivalently defined this process in terms of "node pruning." As a result, for

acyclic $s_{\max}$ graphs, it is easy to see the following.

Proposition 4.2.1. *Let g be an acyclic $s_{\max}$ graph satisfying ordered degree sequence $D = \{d_1, d_2, \ldots, d_n\}$. For $1 \leq k \leq n$, denote by $\tilde{g}_k$ the acyclic graph obtained by removing ("trimming") in order nodes $n, n-1, \ldots, k+1$ from g. Then, $\tilde{g}_k$ is the $s_{\max}$ graph for degree sequence $\tilde{D}_k = \{d_1', d_2', \ldots, d_k'\}$.*

The proof of Proposition 4.2.1 follows directly from our proof of the construction of the $s_{\max}$ graph for trees (see Appendix A). More generally, for graphs exhibiting large $s(g)$-values, properly defined graph operations of link trimming appear to yield simplified graphs with high s-values, thus suggesting a broader notion of self-similarity or invariance under such operations. However, additional work remains to formalize this notion.

4.2.2 Coarse Graining by Collapsing Nodes

A kind of *coarse graining* of a graph can be obtained for producing simpler graphs by collapsing existing nodes into aggregate or super nodes and removing any duplicate links emanating from the new nodes. Consider the case of a tree g having degree sequence $D = \{d_1, d_2, \ldots, d_n\}$ satisfying $d_1 \geq d_2 \geq \cdots \geq d_n$ and connected in a manner such that $s(g) = s_{\max}$. Then, as long as node aggregation proceeds in order with the degree sequence (i.e., aggregate nodes 1 and 2 into $1'$, then aggregate nodes $1'$ and 3 into $1''$, and so on), all intermediate graphs $\tilde{g}$ will also have $s(\tilde{g}) = s_{\max}$. To see this, observe that for trees, when aggregating nodes 1 and 2, we have an abbreviated degree sequence $D' = \{d_1', d_3, \ldots, d_n\}$, where $d_1' = d_1 + d_2 - 2$. Provided that $d_2 \geq 2$ then we are guaranteed to have $d_1' \geq d_3$, and the overall ordering of D' is preserved. Similarly when aggregating nodes $1'$ and 3 we have abbreviated degree sequence $D'' = \{d_1'', d_4, \ldots, d_n\}$, where $d_1'' = d_1 + d_2 + d_3 - 4$. So as long as $d_3 \geq 2$ then $d_1'' \geq d_4$ and ordering of D'' is preserved. And in general, as long as each new node is aggregated in order and satisfies $d_i \geq 2$, then we are guaranteed to maintain an ordered degree sequence. As a result, we have proved the following proposition.

Proposition 4.2.2. *For acyclic $g \in G(D)$ with $s(g) = s_{\max}$, coarse graining according to the above procedure yields smaller graphs $g' \in G(D')$ that are also the $s_{\max}$ graphs of this truncated degree distribution.*

For cyclic graphs, this type of node aggregation operation maintains $s_{\max}$ properties only if the resulting degree sequence remains ordered, i.e., $d_{1'} \geq d_3 \geq d_4$ after the first coarse graining operation and $d_{1''} \geq d_4 \geq d_5$ after the second coarse graining operation, etc.. It is relatively easy to generate cases where arbitrary node aggregation violates this condition and the resulting graph is no longer self-similar in the sense of having a large $s(g)$-value. However, when this condition is satisfied, the resulting simpler graphs seem to satisfy a broader self-similar property. Specifically, for high-$s(g)$ graphs $g \in G(D)$, properly defined graph operations of coarse graining appear to yield simplified graphs in $G(D)$ with high s-values (i.e., such graphs are self-similar or invariant under proper coarse graining), but this has not been proved.

These are, of course, not the only coarse graining, pruning, or merging processes that might be of interest, and for which $s_{\max}$ graphs are preserved, but they are perhaps the simplest to state and prove.

4.2.3 Subgraph-Based Motifs

While graph transformations such as link trimming or node collapse reflect some aspects of what it means for a graph to be self-similar, the graph transformation of random pairwise degree-preserving link rewiring offers additional notions of self-similarity which potentially are even richer and also connected with the claim in the scale-free literature that scale-free graphs are preserved under such rewirings.

For any graph $g \in G(D)$, consider the set of local degree-preserving rewirings of distinct pairs of links. There are $\binom{l}{2} = l(l-1)/2$ pairs of different links on which degree preserving rewiring can occur. Each pair of links defines its own network subgraph, and in the case where g is an acyclic graph (i.e., a tree), these form three distinct types of subgraphs, as shown in figure 4.3. Here rewiring operations that result in nonsimple graphs (shaded) are assumed to revert to the original configuration. Thus defined, rewiring of motif (i) does not result in any new graphs, rewiring of motif (ii) results in one possible new graph, and rewiring of motif (iii) results in two possible new graphs.

Using the notation $d^2 = \sum {d_k}^2$, $s = s(g)$ we can enumerate the number of these sub-

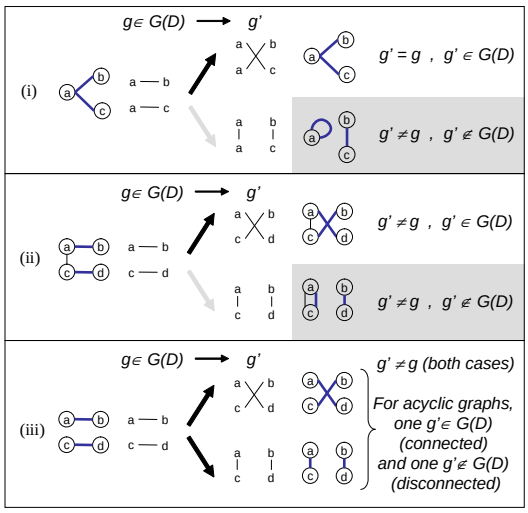

Figure 4.3: Three possible subgraph-based motifs in degree-preserving rewiring in acyclic graphs.

Table 4.1: The numbers of the three motifs and successively the number for each possible rewiring outcome in figure 4.3.

		Outcome from degree-preserving rewiring			
		$g' \in G(D)$ simple connected		$g' \notin G(D)$	
				simple not connected	not simple not connected
Case/Motif	Count	$g' = g$	$g' \neq g$	$g' \neq g$	$g' \neq g$
(i)	$\frac{d^2}{2} - l$	1	0	0	1
(ii)	$s - d^2 + l$	0	1	0	1
(iii)	$\frac{d^2}{2} - s + \frac{l^2-l}{2}$	0	1	1	0
Totals	$\frac{l^2-l}{2}$	$\frac{d^2}{2} - l$	$\frac{l^2+l}{2} - \frac{d^2}{2}$	$\frac{d^2}{2} - s + \frac{l^2-l}{2}$	$s - \frac{d^2}{2}$

graphs as follows:

1. The two links share a common node. There are $\sum_{i=1}^{n} \binom{d_i}{2} = \frac{1}{2}d^2 - l$ possible ways that this can occur.

2. The links have two nodes that are connected by a third link. There are $\sum_{(i,j)\in\mathcal{E}}(d_i - 1)(d_j - 1) = s - d^2 + l$ possible ways that this can occur.

3. The links have end points that do not share any direct connections. There are $\binom{l}{2} - \sum_{i=1}^{n}\binom{d_i}{2} - \sum_{(i,j)\in\mathcal{E}}(d_i - 1)(d_j - 1) = \frac{1}{2}d^2 - s + \frac{1}{2}(l^2 - 2)$ possible ways that this can occur.

Collectively, these three basic subgraphs account for all possible $\binom{l}{2} = l(l-1)/2$ pairs of different links. The subgraphs in cases (i) and (ii) are themselves trees, while the subgraph in case (iii) is not. We will refer to these three cases for subgraphs as "motifs," in the spirit of [86], noting that our notion of subgraph-based motifs is motivated by the operation of random rewiring to be discussed below.

The simplest and most striking feature of the relationship between motifs and $s(g)$ for acyclic graphs is that we can derive formulas for the number of subgraph-based (local) motifs (and the outcomes of rewiring) entirely in terms of d^2, $s = s(g)$, and l. Thus, for example, we can see that graphs having higher d^2 (equivalently higher CV) values have fewer of the second motif. If we fix D, and thus l and d^2, for all graphs of interest, then the only remaining dependence is on s, and graphs with higher $s(g)$-values contain fewer disconnected (case iii) motifs. This can be interpreted as a motif-level connection between $s(g)$ and self-similarity, in that graphs with higher $s(g)$ contain more motifs that are themselves trees, and thus more similar to the graph as a whole. Graphs having lower $s(g)$ have more motifs of type (iii) that are disconnected and thus dissimilar to the graph as a whole. Thus high-$s(g)$ graphs have this "motif self-similarity," low-$s(g)$ graphs have "motif self-dissimilarity" and we can precisely define a measure of this kind of self-similarity and self-dissimilarity as follows.

Definition 4.2.1. *For a graph $g \in G(D)$, another measure of the extent to which g is self-similar is the metric $ss(g)$ defined as the number of motifs (cases i-ii) that are themselves connected graphs. Accordingly, the measure of self-dissimilarity $sd(g)$ is then the number of motifs (case iii) that are disconnected.*

For trees, $ss(g) = s - d^2/2$ and $sd(g) = -s + (l^2 - l + d^2)/2$, so this local motif self-similarity (self-dissimilarity) is essentially equivalent to high-$s(g)$ (low-$s(g)$). As noted previously, network motifs have already been used as a way to study self-similarity and

coarse graining [61, 60]. There, one defines a recursive procedure by which node connectivity patterns become represented as a single node (i.e., a different kind of motif), and it was shown that many important technological and biological networks were self-dissimilar, in the sense coarse-grained counterparts display very different motifs at each level of abstraction. Our notion of motif self-similarity is much simpler, but consistent, in that the Internet has extremely low $s(g)$ and thus minimally self-similar at the motif level. The next question is whether high $s(g)$ is connected with "self-similar" in the sense of being preserved under rewiring.

4.2.4 Degree-Preserving Rewiring

We can also connect $s(g)$ in several ways with the effect that local rewiring has on the global structure of graphs in the set $G(D)$. Recall the above process by which two network links are selected at random for degree-preserving rewiring, and note that when applied to a graph $g \in G(D)$, there are four possible distinguishable outcomes:

1. $g' = g$ with $g' \in G(D)$: the new graph g' is *equal* to the original graph g (and therefore also a simple, connected graph in $G(D)$);

2. $g' \neq g$ with $g' \in G(D)$: the new graph g' is not equal to g, but is still a simple, connected graph in the set $G(D)$ (note that this can include g' which are isomorphic to g);

3. $g' = g$ with $g' \notin G(D)$: the new graph g' is still simple, but is not connected;

4. $g' = g$ with $g' \notin G(D)$: the new graph g' is no longer simple (i.e., it either contains self-loops or parallel links).

There are two possible outcomes from the rewiring of any particular pair of links, as shown in figure 4.3(a) and this yields a total of $2\binom{l}{2} = l(l-1)$ possible outcomes of the rewiring process. In our discussion here, we ignore isomorphisms and assume that all nonequal graphs are different.

We are ultimately interested in retaining within our new definitions the notion that high $s(g)$ graphs are somehow preserved under rewiring provided this is sufficiently random and

degrees are preserved. Scaling is of course trivially preserved by any degree-preserving rewiring, but high $s(g)$ value is not. Again, Figure 4.1 provides a clear example, since successive rewirings can take any of these graphs to any other. More interesting for high $s(g)$ graphs is the effect of *random* rewiring. Consider again the $Perf(g)$ vs. $s(g)$ plane from Figure 4.1. In addition to the five networks from Figure 3.8, we show the $Perf(g)$ and $s(g)$ values for other graphs in $G(D)$ obtained by degree-preserving rewiring from the initial four networks. This is done by selecting uniformly and randomly from the $l(l-1)$ different rewirings of the $l(l-1)/2$ different pairs of links, and restricting rewiring outcomes to elements of $G(D)$ by resetting all disconnected or nonsimple neighbors to equal. Points that match the color of one of the four networks are only one rewiring operation away, while points represented in gray are more than one rewiring operation away.

The connections of the results in Table 4.1 to motif counts is more transparent however than to the consequences of successive rewiring. Nevertheless, we can use the results in Table 4.1 to describe related ways in which low $s(g)$ graphs are "destroyed" by random rewiring. For any graph g, we can enumerate among all possible pairs of links on which degree preserving rewiring can take place and count all those that result in equal or nonequal graphs. In Figure 4.3, we consider the four cases for degree-preserving rewiring of acyclic graphs, and we count the number of ways each can occur. For motifs (i) and (ii), it is possible to check locally for outcomes that produce nonsimple graphs and these cases correspond to the shaded outcomes in figure 4.3. If we a priori exclude all such nonsimple rewirings, then there remain a total of $l(l-1) - s + d^2/2$ simple similar neighbors of a tree. We can define a measure of local rewiring self-dissimilarity for trees as follows. We distinguish between equal, not equal but connected and simple, not connected but simple, and not simple graphs that are similar to each graph with the given motif selected for rewiring. In Table 4.1, the total number of cases (column sum) is $(l^2 - l)/2$, while the total number (row sum) of outcomes is twice that at $l^2 - l$.

Definition 4.2.2. *For a tree $g \in G(D)$, we measure the extent to which g is self-dissimilar under local rewiring by the metric $rsd(g)$ defined as the number of simple similar neighbors that are disconnected graphs.*

For trees, $rsd(g) = sd(g) = -s + (l^2 - l + d^2)/2$, so this local rewiring self-dissimilarity is identical to motif self-dissimilarity and directly related to low $s(g)$ values. This is because only motif (iii) results in simple but not connected similar neighbors.

4.3 The s-metric and Likelihood

While the introduction and exploration of the *s*-metric fits naturally within standard studies of graph theoretic properties, it differs from the scale-free literature in that our structural approach does not depend on a probability model underlying the set of graphs of interest. The purpose of this section is to compare our approach with the more conventional probabilistic and ensemble-based views. For many application domains, including the Internet, there seems to be little motivation to assume networks are samples from an ensemble, and our treatment here will be brief while trying to cover this broad subject. Here again, we show that the $s(g)$ metric is potentially interesting and useful, as it has a direct relationship to notions of graph likelihood. This section also highlights the striking differences in the way that randomness is treated in physics-inspired approaches vs. those shaped by mathematics and engineering.

4.3.1 Probabilistic Approach

The starting point for most probabilistic approaches to the study of graphs is through the definition of an appropriate *statistical ensemble* (see for example [43, Section 4.1]).

Definition 4.3.1. *A statistical ensemble of graphs is defined by*

(i) a set G of graphs g, and

(ii) a rule that associates a real number ("probability") $0 \leq P(g) \leq 1$ with each graph $g \in G$ such that $\sum_{g \in G} P(g) = 1$.

To describe an ensemble of graphs, one can either assign a specific weight to each graph or define some process (i.e., a stochastic generator) which results in a weight. For

example, in one basic model of random graphs, the set G consists of all graphs with a node set $V = \{1, 2, \ldots, n\}$ having l links, and each element in G is assigned the same probability $1/\binom{n}{l}$. In an alternative random graph model, the set G consists of all graphs with node set $V = \{1, 2, \ldots, n\}$ in which the link are chosen independently and with probability $0 < p < 1$. In this case, the probability $P(g)$ depends on the number of links in g and is given by $P(g) = p^l(1-p)^{n-l}$, where l denotes the number of links in $g \in G$.

The use of stochastic construction procedures to assign statistical weights has so dominated the study of graphs that the assumption of an underlying probability model often becomes implicit. For example, consider the four graph construction procedures listed in [43] that are claimed to form *"the basis of network science,"* and include (1) classical random graphs due to Erdös and Renyí [47]; (2) equilibrium random graphs with a given degree distribution such as the *Generalized Random Graph (GRG)* method [35]; (3) "small-world networks" due to Watts and Strogatz [120]; and (4) networks growing under the mechanism of preferential linking due to Barabási and Albert [20] and made precise in [26]. All of these construction mechanisms are inherently *stochastic* and provide a natural means for assigning, at least in principle, probabilities to each element in the corresponding space of realizable graphs. While deterministic (i.e., nonstochastic) construction procedures have been considered [23], their study has been restricted to the treatment of deterministic preferential attachment mechanisms that result in pseudofractal graph structures. Graphs resulting from other types of deterministic constructions are generally ignored in the context of statistical physics-inspired approaches since within the space of all feasible graphs, their likelihood of occurring is typically viewed as vanishingly small.

Using the construction procedure associated with the *general model of random graphs with a given expected degree sequence* considered in [35] (also called the *Generalized Random Graph (GRG) model* for short) we show that the $s(g)$ metric allows for a more familiar ensemble-related interpretation as *(relative) likelihood* with which the graph g is constructed according to the GRG method. To this end, the GRG model is concerned with generating graphs with given *expected* degree sequence $D = \{d_1, \ldots, d_n\}$ for nodes $1, \ldots, n$. The link between nodes i and j is chosen independently with probability p_{ij}, with p_{ij} proportional to the product $d_i d_j$ (i.e., $p_{ij} = \rho d_i d_j$, where ρ is a sufficiently small constant),

and this defines a probability measure P on the space of all simple graphs and thus induces a probability measure on $G(D)$ by conditioning on having degree D. The construction is fairly general and can recover the classic Erdös-Rényi random graphs [47] by taking the expected degree sequence to be $\{pn, pn, \ldots, pn\}$ for constant p. As a result of choosing each link $(i, j) \in \mathcal{E}$ with a probability that is proportional to $d_i d_j$ in the GRG model, different graphs are typically assigned different probabilities under P. This generation method is closely related to the *power-law Random Graph (PLRG)* method [9], which also attempts to replicate a given (power-law) degree sequence. The PLRG method involves forming a set L of nodes containing as many distinct copies of a given node as the degree of that node, choosing a random matching of the elements of L, and applying a mapping of a given matching into an appropriate (multi)graph. It is believed that the PLRG and GRG models are *"basically asymptotically equivalent, subject to bounding error estimates"* [9]. Defining the *likelihood* of a graph $g \in G(D)$ as the logarithm of its probability under the measure P, we can show that the log likelihood (LLH) of a graph $g \in G(D)$, can be computed as

$$LLH(g) \approx \kappa + \rho\, s(g), \tag{4.13}$$

where κ is a constant.

Note that the probability of any graph g under P is given by [102]

$$P(g) \;=\; \prod_{(i,j)\in\mathcal{E}} p_{ij} \prod_{(i,j)\notin\mathcal{E}} (1 - p_{ij}),$$

and using the fact that under the GRG model, we have $p_{ij} = \rho d_i d_j$, where $D = (d_1, \ldots d_n)$ is the given degree sequence, we get

$$\begin{aligned} P(g) &= \rho^l \prod_{i\in\mathcal{V}} d_i^{d_i} \prod_{(i,j)\notin\mathcal{E}} (1 - \rho d_i d_j) \\ &= \rho^l \prod_{i\in\mathcal{V}} d_i^{d_i} \frac{\prod_{i,j\in\mathcal{V}} (1 - \rho d_i d_j)}{\prod_{(i,j)\in\mathcal{E}} (1 - \rho d_i d_j)}. \end{aligned}$$

Taking the log, we obtain

$$\begin{aligned}\log P(g) &= l\log\rho + \sum_{i\in\mathcal{V}} d_i \log d_i + \sum_{i,j\in\mathcal{V}} \log(1-\rho d_i d_j) \\ &\quad - \sum_{(i,j)\in\mathcal{E}} \log(1-\rho d_i d_j).\end{aligned}$$

Defining

$$\kappa = l\log\rho + \sum_{i\in\mathcal{V}} d_i \log d_i + \sum_{i,j\in\mathcal{V}} \log(1-\rho d_i d_j),$$

we observe that κ is constant for fixed degree sequence D. Also recall that $\log(1+a) \approx a$ for $|a| << 1$. Thus, if ρ is sufficiently small so that $p_{ij} = \rho d_i d_j << 1$, we get

$$LLH(g) = \log P(g) \approx \kappa + \sum_{(i,j)\in\mathcal{E}} \rho d_i d_j.$$

This shows that the graph likelihood $LLH(g)$ can be made proportional to $s(g)$ and thus we can interpret $s(g)/s_{\max}$ as *relative likelihood* of $g \in G(D)$, for the $s_{\max}$-graph has the highest likelihood of all graphs in $G(D)$. Choosing $\rho = 1/\sum_{i\in\mathcal{V}} d_i = 1/2l$ in the GRG formulation results in the expectation

$$E(d_i) = \sum_{j=1}^{n} p_{ij} = \sum_{j=1}^{n} \rho d_i d_j = \rho d_i \sum_{j=1}^{n} d_j = d_i.$$

However, this ρ may not have $p_{ij} = \rho d_i d_j << 1$ and can even make $p_{ij} > 1$, particularly in cases when the degree sequence is scaling. Thus ρ must often be chosen much smaller than $\rho = 1/\sum_{i\in\mathcal{V}} d_i = 1/2l$ to ensure that $p_{ij} << 1$ for all nodes i, j. In this case, the "typical" graph resulting from this construction with have degree sequence much less than D, however this sequence will be proportional to the desired degree sequence, $E(d_i) \propto d_i$.

While this GRG construction yields a probability distribution on $G(D)$ by conditioning on having degree sequence D, this is not an efficient, practical method to generate members of $G(D)$, particularly when D is scaling and it is necessary to choose $\rho << 1/2l$. The appeal of the GRG method is that it is easy to analyze and yields probabilities on $G(D)$ with clear interpretations. All elements of $G(D)$ will have nonzero probability with log

likelihood proportional to $s(g)$. But even the s_{max} graph may be extremely unlikely, and thus a naive Monte Carlo scheme using this construction would rarely yield any elements in $G(D)$. There are many conjectures in the scale-free literature that suggest that a wide variety of methods, including random degree-preserving rewiring, produce "essentially the same" ensembles. Thus it may be possible to generate probabilities on $G(D)$ that can both be analyzed theoretically and also provide a practical scheme to generate samples from the resulting ensemble. While we believe this is plausible, it's rigorous resolution is well beyond the scope of this thesis.

4.3.2 Highly Likely Constructions

The interpretation of $s(g)$ as (relative) graph likelihood provides an explicit connection between this structural metric and the extensive literature on random graph models. Since the GRG method is a general means of generating random graphs, we can in principle generate random instances of "scale-free" graphs with a prescribed power-law degree sequence, by using GRG as described above and then conditioning on that degree sequence. (And more efficient, practical schemes may also be possible). In the resulting probability distribution on the space of graphs $G(D)$, high-$s(g)$ graphs with hublike core structure are literally "highly likely" to arise at random, while low-$s(g)$ graphs with their high-degree nodes residing at the graphs' peripheries are "highly unlikely" to result from such stochastic construction procedures.

While graphs resulting from stochastic preferential attachment construction may have a different underlying probability model than GRG-generated graphs, both result in simple graphs having approximate scaling relationships in their degree distributions. One can understand the manner in which high-$s(g)$ graphs are "highly likely" through the use of a simple Monte Carlo simulation experiment. An alternate approach to generating random graphs having a power-law in their distribution of node degree is to use the type of preferential attachment mechanism first outlined in [20] and consider the structural features that are most "likely" among a large number of trials. Here, we generate 100,000 graphs each having 1000 nodes and measure the s-value of each. It is important to note that

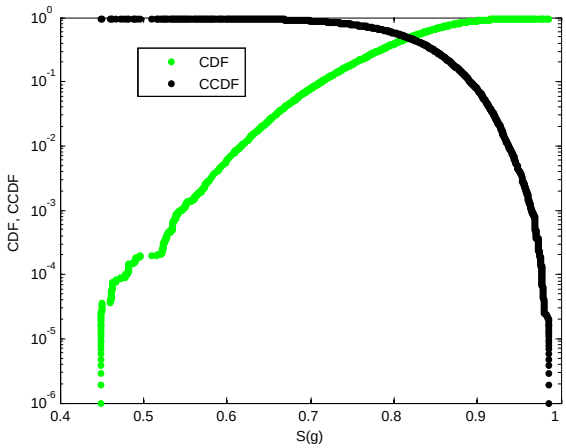

Figure 4.4: Results from Monte Carlo generation of preferential attachment graphs having 1000 nodes. Both the CDF and CCDF are shown.

successive graphs resulting from preferential attachment will have different node degree sequences (one that is undoubtedly different from the degree sequence in figure 3.8(a)), so a raw comparison of $s(g)$ is not appropriate. Instead, we introduce the normalized value $S(g) = (s(g) - s_{min})/(s_{\max} - s_{\min})$ and use it to compare the structure of these graphs. Note that this means also generating the $s_{\max}, s_{\min}$ graph associated with the particular degree sequence for the graph resulting from each trial. Fortunately, the construction procedure in Appendix A makes this straightforward, and so in this manner we obtain the normalized S-values for 100,000 graphs resulting from the same preferential attachment procedure. Plotting the CDF and CCDF of the S-values for these graphs in figure 4.4, we observe a striking picture: all of the graphs resulting from preferential attachment had values of S greater than 0.4, most of the graphs had values $0.5 < S(g) < 0.9$, and a significant number had values $S(g) > 0.9$. In contrast, the graphs in figure 3.8 had values: $S(PA) = 0.52$, $S(HOT) = 0.05$. Again, from the perspective of stochastic construction processes, low-S values typical of HOT constructions are "very unlikely" while high-S values are much more "likely" to occur at random.

With this additional insight into the s-values associated with different graphs, the relationship in the $Perf(g)$ vs. $s(g)$ plot of Figure 4.1 is clearer. Specifically, high-performance networks resulting from a careful design process *are vanishingly rare from a conventional probabilistic graph point of view*. In contrast, the likely outcome of random graph construc-

tions (even carefully handcrafted ones) are networks that have extremely poor performance or lack the desired functionality (e.g., providing connectivity) altogether.

4.4 The s-metric and Assortativity

There is now a growing literature on the importance of correlation structure in networks [49, 95, 94, 41, 92, 111] and how to generate networks having particular correlation structure [67, 110, 32, 78]. A simple measure of correlation structure that has appeared extensively in the literature is the assortativity r which is used to quantify the average tendency of nodes to connect to others having similar degree. It turns out that there is an in inherent relationship between the assortativity and the s-metric, and a closer look at this relationship yields considerable insight into both the diversity within the background set $G(D)$ as well as the interpretation of r itself. In this regard, the assortativity, also noted as Pearson Coefficient, is extremely misleading when measuring graphs with high variability degree sequence since it is directly borrowed from classic graph theory where graphs with low variability graphs dominate.

4.4.1 Assortativity Definition

Recently, Newman [93] introduced the following sample-based measure of graph assortativity as defined by

$$r(g) = \frac{\left[\sum_{(i,j)\in\mathcal{E}} d_i d_j / l\right] - \left[\sum_{(i,j)\in\mathcal{E}} \frac{1}{2}(d_i + d_j)/l\right]^2}{\left[\sum_{(i,j)\in\mathcal{E}} \frac{1}{2}(d_i^2 + d_j^2)/l\right] - \left[\sum_{(i,j)\in\mathcal{E}} \frac{1}{2}(d_i + d_j)/l\right]^2}. \tag{4.14}$$

This relationship can be written as

$$r(g) = \frac{\left[\sum_{(i,j)\in\mathcal{E}} d_i d_j\right] - \left[\sum_{i\in\mathcal{V}} \frac{1}{2} d_i^2\right]^2 / l}{\left[\sum_{i\in\mathcal{V}} \frac{1}{2} d_i^3\right] - \left[\sum_{i\in\mathcal{V}} \frac{1}{2} d_i^2\right]^2 / l}, \tag{4.15}$$

where the first term of the numerator is exactly $s(g)$. Although the assortativity is only a summary statistic for the correlation profile of the graph as a whole, it provides interesting

information nonetheless and is often cited as a key feature distinguishing various classes of complex networks [93, 94, 96, 98].

4.4.2 Measuring Against Background Sets

Here, we argue that $r(g)$ has a natural interpretation as a "centered" and "normalized" version of $s(g)$. In particular, observe that the first term of the denominator in (4.15) is exactly the $s_{\max}$ value within the space $\mathcal{G}(D)$ as defined in (4.10). Accordingly, one can rewrite the assortativity as

$$r(g) = \frac{s(g) - s(g_c)}{s_{\max}^{\mathcal{G}(D)} - s(g_c)}, \tag{4.16}$$

where we refer to g_c as the "center" of the space $\mathcal{G}(D)$.

To see why g_c can be viewed as the center of this space of graphs, we consider the following thought experiment: *what is the structure of a deterministic graph with degree sequence D and having zero assortativity?* In principle, a node in such a graph will connect to any other node in proportion to each nodal degree. In practice, such a graph may not exist for general D, however one can construct a deterministic *pseudograph* $\tilde{g}$ having zero assortativity in the following manner. Let $A = [a_{ij}]$ represent a (directed) node adjacency matrix of nonnegative real values, representing the "link weights" in the pseudograph. That is, links are not constrained to integer values but can exist in nonnegative fractional form. The zero assortative pseudograph will have symmetric weights given by

$$a_{ij} = \left(\frac{d_j}{\sum_{k\in\mathcal{V}} d_k}\right) d_i = \left(\frac{d_i}{\sum_{k\in\mathcal{V}} d_k}\right) d_j = a_{ji}.$$

Thus, the weight a_{ij} for each link emanating out of node i is in proportion to the degree of node j, in a manner that is relative to the sum of all node degrees. In general, the graphs of interest to us are undirected, however here it is notationally convenient to consider the construction of directed graphs. Using these weights, the total weight among all links entering and exiting a particular node i equals

$$\sum_{j\in\mathcal{V}} a_{ij} + \sum_{k\in\mathcal{V}} a_{ki} = d_i + d_i = 2d_i.$$

Accordingly, the total “link weights” in the pseudograph are equal to

$$\sum_{i,j\in\mathcal{V}} a_{ij} = \sum_{j\in\mathcal{V}} d_j = 2l,$$

where again l corresponds to the total number of links in a traditional graph. By extension, the s-metric for the pseudograph $\tilde{g}_A$ represented by connectivity matrix A is calculated as

$$\begin{aligned} s(\tilde{g}_A) &= \sum_{j\in\mathcal{V}} \sum_{i\in\mathcal{V}} \frac{1}{2}\, d_i a_{ij} d_j \\ &= \sum_{j\in\mathcal{V}} \left[\sum_{i\in\mathcal{V}} \frac{1}{2}\, d_i \left(\frac{d_j}{\sum_{k\in\mathcal{V}} d_k} \right) d_i \right] d_j \\ &= \frac{\left(\sum_{j\in\mathcal{V}} d_j^2 \right)\left(\sum_{i\in\mathcal{V}} d_i^2 \right)}{2\left(\sum_{k\in\mathcal{V}} d_k \right)} \\ &= \frac{\left(\sum_{j\in\mathcal{V}} d_j^2 \right)^2}{4l} \\ &= \frac{\left(\sum_{j\in\mathcal{V}} \frac{1}{2} d_j^2 \right)^2}{l}, \end{aligned}$$

showing that $s(\tilde{g}_A) = s(g_c)$. Thus, in terms of its s-value, g_c is equivalent to the center of $\mathcal{G}(D)$.

In principle, one could imagine a deterministic procedure that uses the structural pseudograph $\tilde{g}_A$ to generate the zero assortativity graph among an “unconstrained” background set $\mathcal{G}(D)$. That is, graphs resulting from this procedure could have multiple links between any pair of nodes as well as multiple self-loops and would not necessarily be connected. The challenge in developing such a procedure is to ensure that the resulting graph has degree sequence equal to D, although one can imagine that in the limit of large graphs this becomes less of an issue. By extension, it is not hard to conceive a stochastic process that uses the structural pseudograph $\tilde{g}_A$ to generate a statistical ensemble of graphs having expected assortativity equal to zero. In fact, it is not hard to see why the GRG method [35] is very close to such a procedure.

Note that the total weight in the pseudograph between nodes i and j equals $a_{ij} + a_{ji} =$

$d_i d_j/2l$. As discussed in [70], the GRG method is based on the choice of a probability $p_{ij} = \rho d_i d_j$ of connecting two nodes i and j, and also that in order to ensure that $E(d_i) = d_i$ one needs $\rho = 1/2l$, provided that $\max_{i \neq j \in \mathcal{V}} d_i d_j \leq 2l$. Thus, the GRG method can be viewed as a stochastic procedure that generates real graphs from the pseudograph $\tilde{g}_A$, with the one important difference that the GRG method always results in simple (but not necessarily connected) graphs. Thus, the zero assortativity pseudograph $\tilde{g}_A$ can be interpreted as the "deterministic outcome" of a GRGlike construction method. In fact, it has recently been shown that the statistical ensemble of graphs resulting from the stochastic GRG method has zero assortativity [92].

Thus, the assortativity r (as a summary statistic of graph) captures a fundamental feature of graph structure, one that is closely related to our s-metric. Notice that both centering term and normalization term depend only on D and not on the specific graph, thus, r reflects s is obvious from its definition, but the question is whether a consideration of s by itself provides insight. The calculation of r values for the graphs in figure 3.8 shows that all values are in the interval $[-0.4815, -0.4283]$. In fact, all the simple and connected graphs with the degree sequence as shown in figure 3.8 (a) have r value within $[-0.49, -0.42]$, which makes the assortativity matric almost unable to differentiate any simple connected graph from their ensembles of the same degree sequence. The key reason is that the existing notion of assortativity for an individual graph g is implicitly measured against a background set of graphs $\mathcal{G}(D)$ that is *not* constrained to be either simple or connected. Because r is computed relative to an unconstrained background set, in some cases this normalization (against the unconstrained $s_{\max}$ graph) and centering (against the $\tilde{g}_A$ pseudograph) does a relatively poor job of distinguishing among graphs having the *same* degree sequence, particularly when that degree sequence exhibits high variability. Specifically, one observes that although they have nearly the same assortativity as defined by r, their structural differences are highlighted by s and its normalized values, $s/s_{\max}^{G(D)}$ and $S(g)$, defined as

$$S(g) = \frac{s(g) - s_{\min}}{s_{\max} - s_{\min}}. \tag{4.17}$$

In cases where network performance is measured by the maximum throughput under fixed

node capacities, these structural differences translate to big differences in performance.

4.4.3 Empirical Results

For additional insight into the way in which differences in s translate to differences in r, we extend the previous computational experiment (in section 4.1.4) to values of $r_{\max}$ and $r_{\min}$ within the constrained background set $G(D)$. Note that these values can be computed directly from the corresponding values of $s_{\max}$ and $s_{\min}$. In Figure 4.5(b) we show these values for each of the generated graphs in our experiment. There are several striking features of this plot. The first is that the "normalization" of the s-metric in the calculation of the assortativity r dramatically changes the sense of graph diversity among graphs having a particular D. For values of relatively high $CV(D)$, $r < 0$ and seems largely independent of any diversity as measured by the range in allowable s. In other words, a second important conclusion is that all networks with high $CV(D)$ have $r < 0$ and this seems largely a function of D and not any particular feature of the graph or whether it is a "technological" or "social" network as argued in [98]. This idea has been made previously in [103, 82, 92, 32] and has also been recently argued [122] based largely on empirical observations of real networks having a range of r-values. A third important takeaway is that for small values of $CV(D)$ one observes that small diversity as measured by $s_{\max} - s_{\min}$ translates to a large range of $r_{\max} - r_{\min}$. The last feature we can see is that the differences between the "unconstrained" space $\mathcal{G}(D)$ and the space of simple, connected graphs $G(D)$ may be more important in determining graph properties than other features as measured by aggregate statistics.

It is worth noting that although $r(g) = 1$ is achieved approximately by the s_{max} graph within $\mathcal{G}(D)$ for all graphical D, it is only in very special instances of D where the $s_{\min}$ graph is obtained. Specifically, when $s_{\min} = Z\hat{Z}^T$, then it follows that $r(g) = -1$ if and only if $z_k + \hat{z}_k = z$ (a constant) for each of the k pairs of elements. In other words, although it is true that $r_{\max} = 1$ for arbitrary D, one often observes that $r_{\min} \gg -1$ simply because of the degree sequence D itself.

Based on this analysis, one might reasonably conclude that the assortativity r is not a

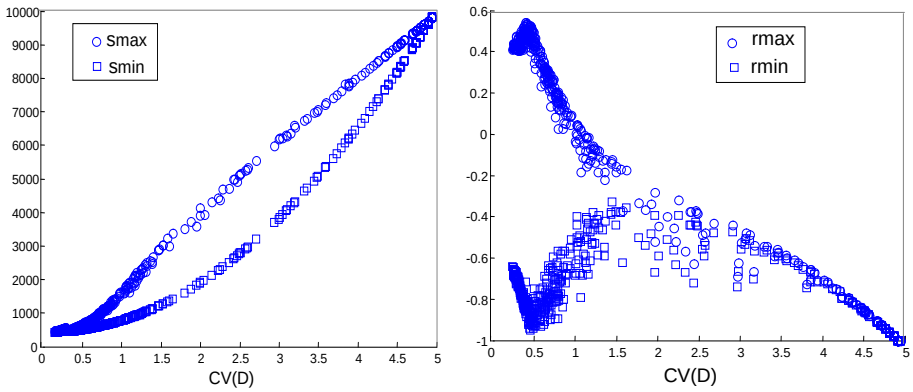

(a) s vs. CV (b) r vs. CV

Figure 4.5: Comparison between s and r with respect to CV

suitable metric for comparing the correlation structure of graphs from different domains. Indeed, it is well understood that a more accurate approach is to consider higher order forms of correlation. Yet the deeper question relates to how one should evaluate any observed correlation structure. Recent efforts by several authors have warned against graph theoretic analysis of networks in isolation. For example, Maslov et al. [81, 82] have argued that a real assessment of a network's correlation structure makes sense only when compared against its "randomized" counterpart. In the context of 'rich-club' ordering in complex networks (i.e., the tendency of high-degree nodes to connect to one another), Colizza et al. [37] have also argued that the presence of high-degree nodes in a given network is enough to ensure that high-degree nodes are connected, and they similarly argue for the need to compare the features of any subject network to a randomized baseline. Thus, important questions include: *What is the appropriate baseline against which to compare graphs?* and *How does this relate to the background set of graphs, as defined by* $G(D)$ *or* $\mathcal{G}(D)$*?*

An inherent challenge in the study of graph diversity is that the combinatorics of even relatively small networks typically result in a space of graphs that is incredibly large. In this study, we have focused on graphs having $n = 100$ (which are about the largest that can be visualized easily) for purposes of exposition, and even here a comprehensive analysis of the elements in $\mathcal{G}(D)$ and $G(D)$ is challenging. In choosing preferential attachment as our primary means for graph generation, we have tried to keep our methods closely tied to the literature so that they may be easily replicated. An alternate approach could have been

to identify specific degree sequences D for which graph isomorphism reduces the number of unique graphs to a small handful and the entire space of graphs (not just $s_{\max}$ and $s_{\min}$) is easily visualized. Identifying and exploring such examples may represent an important step in future work.

The overall message of the results here is that one must carefully consider the inherent diversity of graphs sharing a particular statistical measure when making claims based on any such statistic. Nonetheless, additional work is required to understand fully the way in which graph diversity affects such characterizations. While others have argued for the need to compare against a "randomized" version of the graph, here we have compared against the entire feasible region, as measured by the range $[s_{\min}, s_{\max}]$. The examples here seem to suggest that the distribution of graphs within either $G(D)$ or $\mathcal{G}(D)$ is not uniform, and a general characterization of these distributions is unknown. Ideally, one would like to know more about where the randomized graph sits within the overall space (i.e., is it the "center" of this space?) Moreover, there may be important differences between graph properties that are imposed by structural constraints (e.g., by the degree sequence D) and those relative to what has been randomized.

Although this study provides additional insight into the way in which graph diversity affects one's ability to use aggregate statistics for characterizing complex networks, it has done so primarily for acyclic graphs (i.e., trees), and more work is required to understand the extent to which these same results hold for more general network structures. However, we now present preliminary empirical evidence that suggests the story for nontrees is qualitatively the same.

In Figure 5.4, we show the results of a final experiment in which we again generate trees having $n = 100$ nodes according to attachment rule (4.12) for a range of exponents p. However, to each tree having an initial $l = n - 1$ links we then add an additional kl links by choosing end points probabilistically in correspondence with (4.12). In this manner, we generate graphs having n nodes and a degree sequence D satisfying $\sum_i d_i = 2(k+1)(n-1)$ (i.e., the average degree is $\langle d \rangle \approx 2(k+1)$). Empirical evidence [96] suggests that, for many real networks, $\langle d \rangle < 10$. For each degree sequence D, we then compute the corresponding $s_{\min}$, $s_{\max}$, $r_{\min}$, and $r_{\max}$ values as was done previously. Figure 5.4 shows these values

plotted against the variation of D, represented again as $CV(D)$ and also now normalized as $CV(D)/C_V^{\max}(D)$ for purposes of comparison.

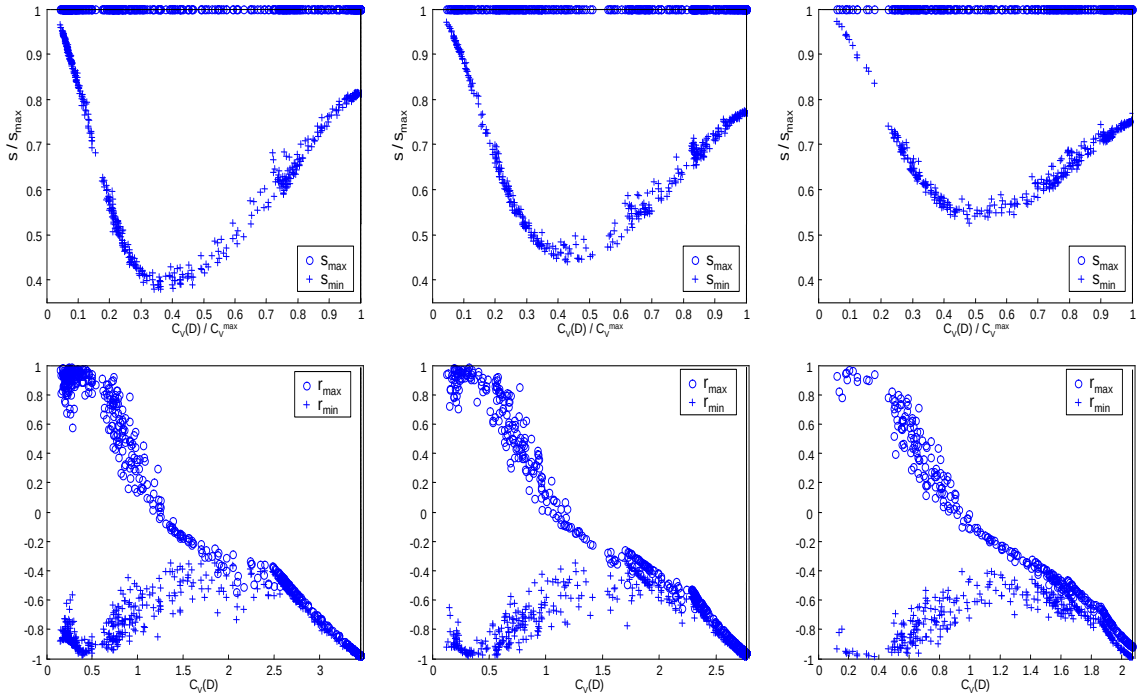

(d) (e) (f)

Figure 4.6: Graph diversity among nontrees. In this experiment, an additional $k(n-1)$ links were added to initial trees of size $n = 100$. (a) $k = 1$, $\langle d \rangle = 3.96$, $C_V^{\max} = 3.4451$. (b) $k = 2$, $\langle d \rangle = 5.94$, $C_V^{\max} = 2.7672$. (c) $k = 4$, $\langle d \rangle = 9.9$, $C_V^{\max} = 2.0701$. In the bottom graphs, variation is measured with $CV(D)$ while in the top graphs it is represented as the normalized $CV(D)/C_V^{\max}(D)$.

One observes for graphs with increasing average degree ($\langle d \rangle \approx 4, 6, 10$ in figure 5.4(a)-(c) respectively) that $CV(D)$ decreases overall but the relative shape of the space of graphs within $G(D)$, as defined by the range $[s_{\min}, s_{\max}]$, remains qualitatively consistent with that of trees. However, the total variation as measured by the distance between $(s_{\max} - s_{\min})/s_{\max}$ decreases with increasing link density. At the same time, for graphs with increasing link density and having degree sequence with $C_V^{\max}(D)$, the difference $s_{\max} - s_{\min}$ is no longer zero in general, indicating inherent diversity even at higher levels of variation. [1]. Graph assortativity as measured by the range $[r_{\min}, r_{\max}]$ is also qualitatively the same as for trees,

[1]However, when the degree sequence D corresponds to a "multistar" (e.g., double-star, triple-star), the overall picture in the upper row of figure 5.4 looks the same, except that the $s_{\min}/s_{\max}$ values jump abruptly to 1 at $C_V^{\max}(D)$, since all "multistars" are isomorphic to one another in $G(D)$.

in that high-$CV(D)$ is enough to dictate that $r < 0$ but that considerable diversity exists for low values of $CV(D)$. Although such results are not conclusive, we view them as generally supportive of graph diversity as we have discussed it here.

4.5 Summary

This chapter provides enhanced understanding towards a theory of scale-free networks by introducing a structural metric, the *s*-metric, that defines one possible measure of the extend to which a graph is scale-free. The *s*-metric is the first one that targets to differentiate between all simple, connected graphs having an identical high variability degree sequence. This structural view has rich and interesting connections to the previously studied graph properties of scale-free networks, such as various notions of self-similarity, likelihood and assortativity. Our approach clarifies much of the confusion surrounding the sensational qualitative claims in the current literature and offers a rigorous and quantitative alternative. We also suggest that when making statements about a graph based on these graph properties one must consider the background set against which these properties are being evaluated.

The functional metric for the Internet topology in the previous chapter together with the structural metric introduced here provide a two-dimensional plane to visualize the diversity of the graph space. In the next chapter, we will talk about the GRAPH of graphs, a connected graph space, which enables a clear understanding of the whole space of graphs from a microscopic viewpoint.

Chapter 5

The GRAPH of graphs

Previously, we project the space of graphs with the same degree sequence into the functional vs. structure plane and derive important relationships among graphs in this extremely diverse space. In this chapter, we propose a new way to view the space of graphs by connecting graphs according to a fundamental microscopic transformation. We call this connected space of graphs the GRAPH of graphs, where each node in the GRAPH represents a graph and a link indicates a local transformation between the two corresponding graphs. The GRAPH of graphs provides a much clearer picture for the whole space of graphs, since we can break this giant space into many subspaces with the same common properties, enumerate all of them, and then explicitly count the number of graphs in each subspace. The GRAPH of graphs also lends perspective on the structural relationship among all the graphs in this domain space. Interestingly enough, many properties of the GRAPH of graphs have direct connections to the properties of graphs inside it. For example, when the GRAPH of graphs represents a domain space of all the simple and connected graphs with the same numbers of nodes and links, the degree of a node in the GRAPH is most relevant to degree variability of the graph that the node represents. While further constrained to the graphs with the same degree sequence, many properties of the GRAPH are related to the graph s-metric.

This chapter is organized as follows. We give motivation and overview of our work in section 5.1. In section 5.2, we first provide graph transformation standard, and propose a fundamental graph transformation method, the general flip, which is proved to that it satisfies the graph transformation standard. In section 5.3, the GRAPH of graphs by the

general flip, G1 is discussed in great detail, including its properties, canonical graphs and subspaces according to graph variability. Then in section 5.4 we propose an extension of the general flip, the degree preserving flip, and the GRAPH of graph (G2) defined by it. We conclude at the end.

5.1 Introduction

The space of graphs is extremely diversity, even when all the graphs share some common large-scale properties, like the same numbers of nodes and links, or degree sequences. Both the functional metric and structural metric introduced in the previous chapters can highlight the differences among graphs having the same degree sequence. Yet, it is not hard to recognize that any macroscopic statistic property is not possible to fully describe a graph since the space of graphs has very high order statistics. In current literature, many graphic metrics are proposed to measure the properties of these graphs, however there is no systematic work to evaluate the effectiveness of these graphic metrics. Some metrics which are used to investigate the properties of one network may not be informative for the other networks because different networks exist for different purposes. For example, the performance defined as the maximum throughput for the Internet may not make any sense for the metabolic network. Many metrics may be misleading, especially when they are constructed against a certain background set. As pointed out in section 4.4 graph assortativity r is implicitly measured against a background set of unconstrained graphs, leading to a substantial bias when used as a metric to differentiate constrained graphs with highly variable degree sequences. For these well evaluated metrics, exhaustively enumerating and calculating their value may not be feasible due to intensive computational requirements. If many metrics of two graphs have similar values, should we conclude that these graphs are essentially the same, or we have to explore one more metric?

While more work needs to be done to refine these macroscopic properties in order to thoroughly evaluate different graphs in the graph space, in this chapter, we propose an alternative approach to study the space of graphs by introducing a very fundamental relationship among these graphs according to their microscopic structural similarities and

differences. This relationship is called the general flip, a way to transform one graph to another by changing only one link in the graph. We can establish a connection between any two graphs if they can be transformed to the other by one step general flip. In contrast to the current literature on studying each individual graph model for complex networks where each graph is an isolated point, our work focus on the entire space of graphs as a connected network. As studying the interconnections among network components is an important step to understand complex networks, building bridges among these isolated graphs will provide deep insight into the space of graphs as well as the fundamental similarities and differences of these graphs. When relationships zoom into microscopic structural level, it gives a clear picture of how each graph can be transformed to another and how different two graphs are.

The domain space of graphs we focus on is the set of all the simple and connected graphs with the same numbers of nodes and links, where the set of all graphs with the same degree sequence is a subspace. The general flip is the most fundamental transformation in this domain space and any other transformations in it can be performed as several steps of general flips. For example, we later will introduce the degree preserving flip, which is a special case of two steps of general flip and can maintain the degree sequence of the graphs. We will also show that the general flip can reach all the graphs in the domain space, that is the general flip satisfies the generality defined by [79].

With the aid of the general flip, the space of graphs is no longer a set of isolated graphs, but a *graph* of graphs (we call G1) where each node is a graph and each link represents a local general flip between the two corresponding graphs. G1 contains all the simple and connected graphs with the same numbers of nodes and links and it can be much more complicated than graphs within it since the size of G1 usually exponentially increases with the size of graphs. However a careful study of the properties of the GRAPH of graphs provides many interesting results relating to the properties of graphs, such as the variability of degree sequence and the *s*-metric of the graphs. For example, we prove that the degree of each node in G1 is proportional to the variability of the corresponding graph degree sequence, therefore graphs with the largest number of neighbors are those with the highest variability in their degree sequence. We also propose a heuristic way to construct a graph

with the highest variability, which is called the canonical graph, and prove that all graphs can be transformed to this canonical form, therefore G1 is connected. The canonical form is also the graph that is most likely to appear when we take a random walk on G1. All these indicate in the space of graphs with the same numbers of nodes and links, variability of the degree sequence plays an important role to measure the graphs. Furthermore by breaking this giant space of graphs into many countable subspaces and then enumerating all the graphs in each subspace, we can obtain a much clearer picture of the graph space which has never been understood thoroughly. In fact, the graphs with lower degree variability completely dominate the whole space. Even though each high degree variability graph has higher probability in random walk, it is much more likely to arrive at low degree variability ones due to this reason.

As an extension of the general flip, we introduce the degree-preserving flip, which is a special case of two step general flips and it can also keep the same degree sequence for the graph. The connected graph space defined by degree-preserving flip is called G2 in which all the graphs are simple and connected with a fixed degree sequence, therefore a subspace of G1. The important discovery of this space is that the degree of each node in G2 is directly related to the *s*-value of the graph defined before. Scale-free graphs (i.e., high *s*-value graphs) have more neighbors and are more likely to appear than other graphs when a random walk is performed on G2. These suggest that when the degree sequence is fixed, the *s*-metric is crucial to differentiate these graphs, which is consistent with our previous argument. Unlike G1, G2 can be disconnected, which means that the degree-preserving flip cannot transform one graph to another in some graph spaces. However, we prove that when any of the graph in G2 has a diameter greater than or equal to three, G2 is connected.

5.2 The General Flip

Before getting into the details of general flip, we introduce some basic standards for graph transformations.

5.2.1 Transformation Standards

A graph transformation may be arbitrarily defined, however, there are some basic properties or standards that the transformation should satisfy as below (originally defined in [79]):

- **Soundness**: No transformation maps to graphs which are not in the domain space. Here we study the domain space of all the simple and connected graphs, and restrict to the graphs with the same numbers of nodes and links for the general flip, and we further constrain the domain to the graphs with the same degree sequence for the degree-preserving flip.
- **Generality**: The transformation process does not converge to a specific graph. All graphs can be reached by this transformation and the probability for each graph to be arrived should be nonzero at limit.
- **Feasibility**: The transformation can be described by a simple (distributed) routine changing only a small number of edges of the graph, so that it can be easily implemented.

We start by repeating some notations of graph theory as defined in previous chapters. A graph is defined by a finite node set $\mathcal{V} = \{1, 2, 3, \ldots, n\}$ of size $n = |\mathcal{V}|$ and a link set $\mathcal{E} := \{(u, v) : u, v \in \mathcal{V}, u \neq v\}$. Let d_i denote the degree (i.e., number of connections) of node i, and call $D = \{d_1, d_2, \ldots, d_n\}$ the degree sequence of the graph. Denote the number of links as $l = |\mathcal{E}|$ and $\sum_{i=1}^{n} d_i = 2l$.

5.2.2 The General Flip

The general flip is defined as follows: consider node $v \in \mathcal{V}$ and its two neighbors u and w ($u, w \in \mathcal{V}$), such that u, w are not connected, change the link (u, v) to (u, w), or change the link (v, w) to (u, w) (see figure 5.1). When u, v, w are fully connected to each other, no transform happens and the graph remains the same. The domain space of graphs by this transformation is all the simple and connected graphs with the same numbers of nodes and links.

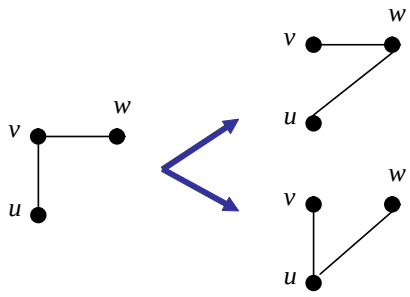

Figure 5.1: The general flip.

Notice that during flipping, among the two links considered, only one link is changed and the other link remains the same. This unchanged link is crucial for maintaining the connectivity of the whole graph. Although there might be a lot of ways to change one graph to another while preserving the number of nodes and links in the graph, the general flip we defined makes it easy to preserve the simplicity and connectivity of graphs since it only requires local information of three nodes. Methods such as random graph rewiring [118] have to periodically check the connectivity of the graph in order to keep it connected, which requires the globe information of the graph structure, therefore has much higher computational complexity. The general flip is also the most fundamental transformation methods for the space of simple and connected graphs having the same numbers of nodes and links, and we will prove that any other method can be performed by general flip within finite steps in this domain space.

We can prove the general flip satisfies the basic standards of graph transformation defined above.

Lemma 5.2.1. *The general flip is sound.*

Proof. The domain space of the general flip is all the simple and connected graphs with the same numbers of links and nodes. The general flip does not add or remove any node and link, therefore it remains the same numbers of nodes and links in the graph. The general flip also keeps the connectness of the graph since the new link added to the graph makes the three nodes reconnected to each other, thus the whole graph stays connected. Furthermore, if the original graph is simple, the transformed graph is also simple since no self-loop or

duplicated link is introduced. The general flip keeps graphs in the same domain space therefore it is sound. □

Lemma 5.2.2. *The general flip is feasible.*

Since the flip only requires the local information of three nodes and it can be described by a simple (distributed) routine by changing only 2 connected links in the graph, the general flip maintains the feasibility.

Lemma 5.2.3. *The general flip is general.*

Proof. We can prove the generality by first introducing a canonical graph in the domain space and then proposing a series of general flips such that all the graphs in the domain space can be transformed into this canonical graph. Realize that the general flip is reversible, therefore all the other graphs can be reached from the canonical graph.

Definition 5.2.1. *The canonical graph for general flips consists of a set of star nodes and a set of edge nodes. All the star nodes are fully connected with each other, and all star nodes except for one are connected to all edge nodes. There are no extra links among edge nodes.*

In the canonical graph, assume the number of star nodes as x, then there will be $x-1$ of nodes having degrees $n-1$ and one special star node having degree $n-1-y$ where $0 \leq y \leq n-1-x$. Correspondingly, y edge nodes having degrees $x-1$ are only connected to the $x-1$ nonspecial star nodes, and $n-x-y$ edge nodes having degrees x which are connected to all the star nodes. The canonical graph is unique (except for its isomorphic graphs) which means that when fixing n and l, x and y is unique. For example, when increasing x, i.e., changing some edge nodes to star nodes, there will be more than one star node that cannot be fully connected to all the other nodes, while decreasing x, extra links has to be put among edge nodes. There is a special case in which $y = 0$ and any node with degree x can be either counted either as a star node or as an edge node, but the graph is the same, and we count this as an edge node without loss of generality.

Now we can define a series of general flips from any graph to the canonical form as follows. Choose one node u_1 as the first star node. From any other node $v \neq u_1$, find a shortest path to u. If the length p of the shortest path is greater than one (i.e., u_1, v are not

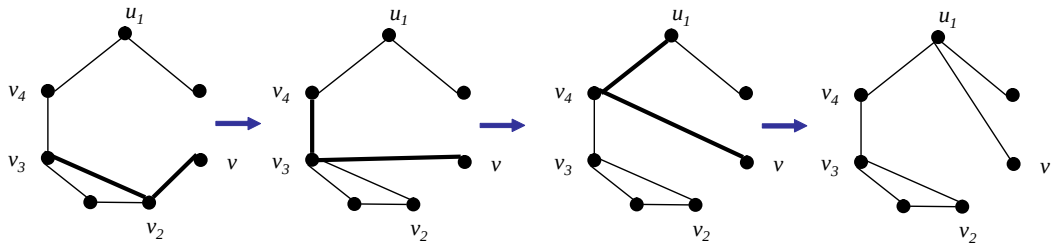

Figure 5.2: A series of general flips to make node v connect to u_1.

directly connected), without loss of generality, we assume the path as $(v, v_2, v_3, \ldots, v_p, u_1)$. Since it is the shortest path, all the nodes on the path form a chain, that is, one node cannot connect to any other nodes on the path except for its two nearest neighbors on the path. Therefore we can perform a series of general flips as: flip $(v, v_2), (v_2, v_3)$ to $(v, v_3), (v_2, v_3)$, then flip $(v, v_3), (v_3, v_4)$ to $(v, v_3), (v_2, v_3)$, and so on, till flip $(v, v_p), (v_p, u_1)$ to $(v, u_1), (v_p, u_1)$ (see figure 5.2). For all the other nodes which are not connected to u_1, perform similar general flips until u_1 is connected to all of them. The general flip can be blocked by a triangle, however, along the shortest path, there is no such triangle that can block the flip transformation. Otherwise we can find a shorter path by going through the shortcut formed by the third link in the triangle which is not on the shortest path.

For a tree, after a series of transformations, the resulting graph is a star in which the degree of the star node is $n-1$ and all the other nodes have degree 1. A star is the canonical form of all trees and the process terminates. In nontree cases, the current graph will be starlike and there is one node connecting to all the other nodes (edge nodes), while edge nodes will have some extra connections among them. Now pick one of the edge nodes as the second star node u_2, for any other edge node v with degree greater than one and not connected to u_2, do the following series of flips: flip $(v, u_1), (u_1, u_2)$ to $(v, u_2), (u_1, u_2)$, and pick any neighbor of v other than u_2, denoted as v_1, flip $(v, v_1), (v_1, u_1)$ to $(v, u_1), (v_1, u_1)$ (see Figure 5.3 for details). Here v_1 exists since the degree of v is greater than one, and v_1 is connected to u_1 since the flip process always tries to maintain the connections between u_1 and any other nodes, even if they could be changed in the intermediate process. By this process, we can make all the nodes with degrees greater than or equal to two to connect to the second star node.

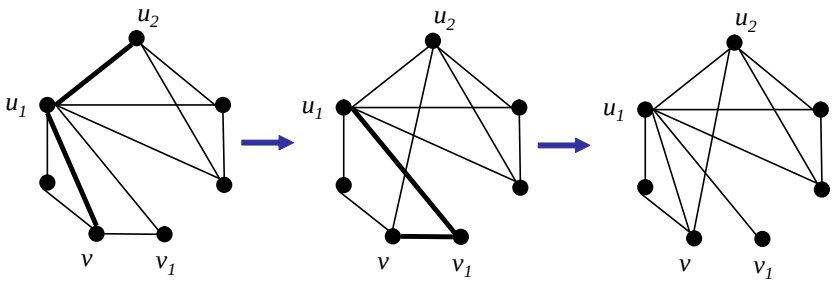

Figure 5.3: A series of general flips to make any edge node v with degrees greater than one to connect to the second star node u_2.

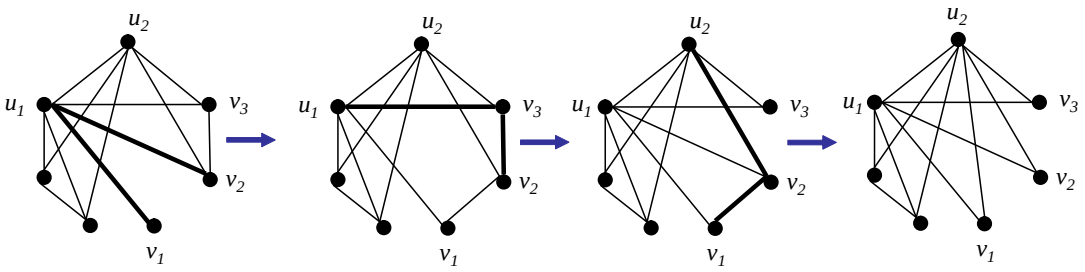

Figure 5.4: The balancing process to add link (v_1, u_2) and remove link (v_2, v_3)

After these transformations, we obtain a double star, where there are two star nodes connecting all the edge nodes except for some possible degree one edge nodes only connecting to the first star u_1. In this double star, if there coexist degree one edge nodes and the edge nodes with degrees greater than two, we call this graph *not balanced*. We introduce a process to balance the double star. Suppose v_1 is one of the nodes with degrees equal to one, and v_2 and v_3 are nodes with degrees greater than two and there is a link between v_2 and v_3. The balancing process is defined as the following flip process: flip $(v_2, u_1), (v_1, u_1)$ into $(v_2, v_1), (v_1, u_1)$, flip $(v_2, v_3), (v_3, u_1)$ into $(v_2, u_1), (v_3, u_1)$, flip $(v_1, v_2), (v_2, u_2)$ into $(v_1, u_2), (v_2, u_2)$ (see figure 5.4). It is easy to show the existence of these links and that these general flips cannot be blocked. After this balancing process, we get a double star in which all the edge nodes have degrees either less than or equal to two, or all the edge nodes have degrees greater than or equal to two. In the first case, the flip process terminates and we obtain the canonical form, while for the latter case, we change one edge node to a star node, and continue.

In summary, the canonicalization process includes two sub-processes, one is called the *centralization process*, in which we connect the edge nodes to the newly added star node, if the edge nodes have degrees greater than x, where x is the number of current star nodes. The centralization process is followed by the *balancing process*, in which we balance the links among edge nodes so that when the graph contains x star nodes, either all the edge nodes have degrees less than or equal to x, or all the edge nodes have degrees greater than or equal to x[1]. In the former case, the process terminates and we get the canonical graph, while in the latter case, we move one edge node to the set to star nodes and continue with the centralization process. In the final canonical graph, the number of star nodes x is uniquely determined by the numbers of nodes and links.

Since any graph can be transformed to the canonical form and each general flip process is reversible, this means that the canonical form can also be changed to any other graph, and as a result all the graphs are reachable. As a result, the final stationary distribution for any graph is nonzero according to [45]. We will study the stationary distribution of each graph later. This finishes the proof of the generality of the general flip. □

Notice that we can define the canonical graphs arbitrarily since any graph can be transformed to others by general flips. However, the canonical graph we define here has many special properties to be discussed later. Also we point out there are many isomorphic canonical graphs since we can pick the star node arbitrarily each time.

In the domain space of simple and connected graphs with the same numbers of links and nodes, the general flip can be considered as the most fundamental transformation method in this space. Since any two graphs can be interchanged to each other by performing a series of general flips, any other transformation method within this space can be achieved by a series of general flips. For example, the degree-preserving flip we will introduce later is a special case of two step general flips.

[1] In the case when all the edge nodes have degrees equal to x, we count it as the first case.

5.3 G1

The graph transformation defines a connection between any two graphs, therefore the space of graphs is now a *GRAPH* of graphs which consists not only nodes representing all the graphs, but also links describing graph transformations. We define G1 as the GRAPH of graphs according to the general flip, such that each node in G1 represents a simple and connected graph with the same numbers of nodes and links as the other graphs in G1, and two nodes share a link if and only if the underlying graphs represented by the nodes can be transformed to each other by one step general flip. Here we reuse the notation in the previous section and define g as a simple and connected graph, and we also denote g as the node in G1 that represents the graph g.

5.3.1 Properties of G1

G1 can be quite complicated due to the huge number of graphs in the space. However, exploring G1 reveals many interesting properties and provide us a much clearer picture of the space of graphs. One immediate property can be derived is that G1 is connected because general flip is general as shown in the previous section. Moreover, we can roughly estimate the number of graphs in G1 and the diameter of G1. Most interestingly, we find that the degree of each node g in G1 is proportional to the variability of the degree sequence of graph g that this node represents.

Lemma 5.3.1. *The diameter of G1 is bounded by the order of* n^2.

The number of nodes is equal to the total number of all the simple and connected graphs with the same numbers of nodes and links. This number is huge and it can exponentially increase with n in general. For example, if the graph is acyclic, i.e. the number of links $l = n - 1$, the total number of trees is n^{n-2} according to [117]. If each graph in G1 has n nodes and l links, without considering whether it is connected or not, the total number of possible graphs in G1 is $M!/l!$, where $M = \binom{n}{2}$ is the total number of slots for l possible links. This number is greater than $l^{M-l} = l^{n(n-1)/2-l}$ and can exponentially increases with n. Due to the existence of the giant component when the number of links l is big, we can

expect the number of connected graphs is not significantly smaller than this estimation. However, the diameter of G1 is polynomial in n, i.e. the logarithmic of the size of G1. We can derive this by counting the total number of general flips it takes from any graph to the canonical form.

Proof. Recall that the transformation from any graph to the canonical form includes a set of centralization processes and a set of balancing processes. In the first centralization process, i.e. when picking the first star node, the maximum number of flips for any node to connect to the star node is the shortest distance from this node to the star node, therefore it is bounded by D, the diameter of the graph, and the first centralization process will take less than Dn flips. After that, each later centralization process takes less than $2n$ flips since the diameter of the graph changes to two after the first centralization process. In the balancing process, each balance step includes three flips and the total flips for each balancing process is less than $3n$. The total number of star nodes depends on the ratio between the numbers of links and nodes, which is approximated to the rounded integer of $n - \sqrt{n^2 - 2l}$. Adding all these flips together, the total number of flips from any arbitrary graph to the canonical form is about $Dn + 3n + 5n(n - \sqrt{n^2 - 2l})$ flips, which is bounded by the order of n^2. This means from any node in G1, we can always find a path to the node which represents the canonical graph within a distance on the order of n^2. As a result, the diameter of the G1, i.e., the maximum shortest distance between any two nodes, is less than or equal to the twice of the distance from any node to the canonical node, and therefore it is also bounded by the order of n^2. □

Lemma 5.3.2. *The degree of each node g in G1 equals $\sum_{i=1}^{n} d_i^2 - 2l - 6\Delta$, where Δ is the total number of triangles in the graph g, d_i and l represent the degree and number of links in the graph g.*

Proof. The degree of each node in G1 is determined by the number of one-step general flips that can happen in the graph represented by the node. This is related to the number of 2-motifs in the graph and the number of triangles in the graphs. A 2-motif is defined in section 4.2 and for any graph with degree sequence $d_1, d_2, \ldots, d_n$, the total number of

2-motifs is

$$\sum_{i=1}^{n} \binom{d_i}{2} = \frac{1}{2} \sum_{i=1}^{n} d_i^2 - l.$$

If the 2-motif chosen to perform the general flip also forms a triangle, the general flip cannot happen, therefore we must remove this case when counting the number of total general flips. Any triangle in the graph is included in three 2-motifs, and we denote the number of triangles in the graphs as Δ. If a 2-motifs does not contain a triangle, it in fact can perform two different general flips, which corresponds two different neighbors that the graph can reach in G1. Therefore degree of each node in G1 is calculated as:

$$\sum_{i=1}^{n} d_i^2 - 2l - 6\Delta. \tag{5.1}$$

□

As introduced in section 4.1.4, $\sum_{i=1}^{n} d_i^2$ is directly related to the *CV*, a measure of the variability of a graph degree sequence, when the numbers of nodes and links are fixed. Therefore, the graph with the higher variability degree sequence will have more neighbors than the graph with lower variability degree sequence. If a graph has more triangles, which corresponds to a higher clustering coefficient [120], it has fewer neighbors than those with lower clustering coefficients.

Lemma 5.3.3. *When performing a random walk from one node on G1 with equal probabilities to all its neighbors, the final stationary distribution for staying at that node is proportional to* $\sum d_i^2 - 2l - 6\Delta$.

As before, d_i, l, Δ are the parameters of the graph that represented by the node. When we do a random walk on a graph, from any node u, jump to a uniformly randomly chosen neighbor, the stationary distribution of any node u with degree $d(u)$ is equal to $d(u)/\sum_u d(u)$ according to [45]. Since the degree of a node in G1 is equal to $\sum d_i^2 - 2l - 6\Delta$, the final stationary distribution of staying at any graph is proportional to $\sum d_i^2 - 2l - 6\Delta$. Therefore, a graph with higher variability degree sequence is easier to appear than a graph with lower variability degree sequence during the random walk.

5.3.2 The Canonical Graph in G1

The canonical graph (and its isomorphic graphs) of the general flip was introduced in the previous section to prove that all the other graphs can be transformed into this form within finite steps. In fact, the canonical graph is the most special graph in G1.

Lemma 5.3.4. *The canonical graph g is the one with the highest CV, therefore corresponds to the node g with the highest degree and the highest stationary distribution in G1.*

Any optimization problem related to graphs can be NP hard due to connectivity and integer constraints. Fortunately, we can prove that a simple variation of constructing canonical graph can achieve the global maximal of the CV. Starting from any graph, since the star nodes can be chosen arbitrarily and the final canonical graphs are just isomorphic to each other, we can always pick the highest degree node among all the remaining edge nodes as the new star node and connect it to all the other edge nodes, while maintaining the degrees of the previous star nodes. Whenever finishing one connection from any edge node to the star node, the degree of the star node will increase by 1 therefore it will remain the highest degree among all the edge nodes. Assume before connecting edge node v to the newly added star node x, the star node has degree d_x. The essence of centralization process for each edge node v is to break the connection between v to any of its edge node neighbor u and connect v to the star node x. In this process, degree of v remains the same, and degree of x will be increased by one at the expense of decreasing the degree of u by one, while all the other nodes remain the same degrees. The change of CV can be calculated as:

$$(d_x + 1)^2 + (d_u - 1)^2 - d_x^2 - d_u^2 = 2(d_x - d_u + 1).$$

Since star node has a degree higher than or equal to the degree of any edge nodes, i.e., $d_x \geq d_u$, the CV of the graph strictly increases when any edge node is connected to a star node. Each centralization process for one star node consists of procedures connecting all the edge nodes to the star node, therefore each centralization process strictly increases the CV.

However, the intermediate balancing process may not always increase the CV and

whether the *CV* can be increased or not depends on the degree of edge nodes with extra links.

However, we can introduce a small variation of the canonicalization process where we do not perform the balancing process immediately after each centralization process, instead performing it only after the centralization process for the *x*th star node (*x* is the total number of star nodes in the final canonical graph). With this variation, we are able to directly prove the global maximization of the *CV* for the canonical graph. Notice that *x* and *y* (the number of edge nodes with degree $x-1$ as defined before) should be precalculated according to the previous two interleaving processes. From the previous proof, we can see that the centralization process is strictly increasing *CV* regardless whether the balancing process is performed or not, as long as each time the new star node picked is the highest degree nodes among all the edge nodes. Assume after *x* centralization processes without any balancing process, the degrees of the *x* star nodes are $n-1-a_1, n-1-a_2, \ldots, n-1-a_x$, where $a_i \geq 0$, $1 \leq i \leq x$. The edge nodes are divided into two parts according to their degree. For these with degree greater than *x*, we assume their degree be $x+e_1, x+e_2, \ldots, x+e_k$ ($e_i > 0$, $1 \leq i \leq k$, $0 < k < n-x$ [2]), and for those with degree less than or equal to *x*, we assume their degree be $x-o_1, x-o_2, \ldots, x-o_h$ ($o_i \geq 0$, $1 \leq i\ h$, $h = n-x-k$. In the former case, these edge nodes are connected to all the star nodes and there are some extra links among the edge nodes, while in the latter case, each edge node *i* is connected to $x-o_i$ star nodes and there are no extra links. The essence of the balancing process is to rearrange the extra links among higher degree edge nodes to connect lower degree edge nodes to star nodes. The *CV* of the graph before rearrangement is

$$\begin{aligned} & \sum_{i=1}^{x}(n-1-a_i)^2 + \sum_{i=1}^{k}(x+e_i)^2 + \sum_{i=1}^{h}(x-o_i)^2 \\ = \quad & x(n-1)^2 - 2(n-1)\sum_{i=1}^{x}a_i + \sum_{i=1}^{x}a_i^2 + (n-x)x^2 + 2x\sum_{i=1}^{k}e_i + \sum_{i=1}^{k}e_i^2 - 2x\sum_{i=1}^{h}o_i + \sum_{i=1}^{h}o_i^2. \end{aligned} \tag{5.2}$$

[2] When $k = 0$, there is no need to do balancing process and each process strictly increases *CV*.

After the balancing process, *CV* changes to

$$(x-1)(n-1)^2 + (n-1-y)^2 + y(x-1)^2 + (n-x-y)x^2.$$

Assume $y < h$, therefore y of the lower degree edge nodes increase their degree to $x-1$, and $h-y$ of them changes their degree to x. The degree of all the higher degree edge nodes drops to x, therefore all the links among them are extra links which will be moved to the slots between lower degree edge nodes and star nodes. Assume the number of extra links among higher degree edge nodes to be L, and $L = \frac{1}{2}\sum_{i=1}^{k} e_i$. After obtaining these L links, the lower degree edge nodes will increase degree to x ($h-y$ of them) or $x-1$ (y of them), therefore $\sum_{i=1}^{h} o_i = L + y$. Similarly, all but one star nodes will increase degree to $n-1$, therefore $\sum_{i=1}^{x} a_i = L + y$. Plug these into equation (5.2), we can get the change of the *CV* after balancing processing is

$$y^2 + y + 2(n-1-x)L - \left(\sum_{i=1}^{x} a_i^2 + \sum_{i=1}^{h} o_i^2 + \sum_{i=1}^{k} e_i^2\right). \tag{5.3}$$

To check whether this term is always greater than or equal to zero, we should look at the worst case scenario where $\left(\sum_{i=1}^{x} a_i^2 + \sum_{i=1}^{h} o_i^2 + \sum_{i=1}^{k} e_i^2\right)$ is maximized. The constraint here are $\sum_{i=1}^{k} e_i^2 = 2L$, $\sum_{i=1}^{x} a_i^2 = L + y$, $\sum_{i=1}^{h} o_i^2 = L + y$, $a_i \geq, e_i \geq, o_i \geq 0$. Of course, one implicit constraint is that the graphs before and after balancing process should be simple and connected, which means we cannot arbitrarily set these values. In fact, this problem is equivalent to solving two separated questions. One is given L links, how to put them among the k nodes with degree x, such that the current degree of these nodes $x + e_i$ maximizes $\sum_{i=1}^{k} e_i$. This can be thought as that assume each node have degree 0, and how to put L links among these to maximize the *CV*. The only difference is that we do not require these L links fully connect to the k nodes, since these k nodes connect to the whole graph already. Maximum $\sum_{i=1}^{k} e_i$ is achieved when one node get L connections and the rest L of them get one connection, therefore

$$\max \sum_{i=1}^{k} e_i = L^2 + L. \tag{5.4}$$

The other problem can be thought in a similar way, but it requires put the L links between the star nodes and the lower degree edge nodes. We can achieve the maximal of $\sum_{i=1}^{x} a_i^2 + \sum_{i=1}^{h} o_i^2$ when one star node is connected to $L + y$ lower degree edge nodes, or when one edge node is connected to $L + y$ star node. In either case

$$\sum_{i=1}^{x} a_i^2 + \sum_{i=1}^{h} o_i^2 = (L + y)^2 + (L + y). \tag{5.5}$$

Equations (5.4) and (5.5) describe the minimum increment of CV by the balancing process, and plug these into equation (5.3), we can get that this amount of change equals to $2(n - x - L - 2)L$ which is always greater than or equal to zero since the number of total nodes n at least contains x star nodes, $L + 1$ higher degree edge nodes and 1 lower degree edge nodes if $L > 0$. This means that the balancing process always increases CV (in the case when $L = 0$, there is no need to do the balancing process).

In summary, if the canonical graph is obtained by performing x centralization processes (where x is the number of star nodes in the final canonical graph) and a balancing process, we can prove that the CV of the canonical graph achieves the global maximal among all the simple and connected graphs having the same numbers of nodes and links, since all these processes monotonically increase CV and we achieve a unique CV (despite of graph isomorphism) no matter whatever the initial graph is.

5.3.3 Exploring G1

We have shown that the degree variability CV is an important measure in G1 and the canonical graph we defined has the highest CV, therefore it has the most number of neighbors and each individual canonical graph has higher probability to appear than other individual graph if we take random walk on G1. However, the random walk may not more frequently arrive at the canonical graphs, if the number of canonical graphs is significantly smaller than these of the lower CV graphs. In this section, we explore more details of the space of simple and connected graphs having the same numbers of links and nodes. More specifically, we break G1 into many countable subspaces so that we can explicitly derive

the total number of graphs according to their degree sequences, and draw the relationship between the final stationary distribution of graphs with respect to their *CV*. Without surprise, low *CV* graphs dominate the space of graphs and we can conclude the high *CV* graphs cannot be generated randomly, but from specific design.

To get know more details on the spaces of graphs, we can divide G1 into many subspaces and further divide each subspace into even smaller subspaces till we can explicitly count the number of graphs in that subspace. The space of graphs containing the same numbers of nodes and links can be partitioned into many subspaces, each containing all the graphs with the same unlabelled degree sequence. Two degree sequences are considered having the same unlabelled degree sequence if they are the same when they are ordered. For each subspace of graphs with the same unlabelled degree sequence, we can divide it into smaller subspaces each of which has the same labelled degree sequence. For each labelled degree sequence, we can exactly calculate the number of labelled graphs for acyclic graphs. Here we use acyclic graphs as an example to illustrate the relationship between the variability of a degree sequence and the total number of labelled graphs with that degree sequence.

Given the space of graphs with n nodes and $n-1$ links, all the possible unlabelled degree sequences can be derived by enumerating all the combinations. In fact, this problem is exactly the same as an integer partition problem without constraint. An unconstrained integer partition problem for number m is stated as follows [57]:

$$m = x_1 + x_2 + \cdots + x_m, \qquad x_1 \geq x_2 \geq \cdots \geq x_m$$

For example, when $m = 4$, the unconstrained integer partition problem is stated as:

$$\begin{aligned} 4 &= 4+0+0+0, \\ 4 &= 3+1+0+0, \\ 4 &= 2+2+0+0, \\ 4 &= 2+1+1+0, \\ 4 &= 1+1+1+1 \end{aligned}$$

We obtain all the sequences $(4,0,0,0),\ldots,(1,1,1,1)$ and add two zeros such that the total length of each sequence is six. If adding 1 to each number in each sequence, the sequences we obtain are exactly the same as all the degree sequences of trees when the number of node equals $n = m + 2 = 6$, that is

$$(5,1,1,1,1,1)$$

$$(4,2,1,1,1,1)$$

$$(3,3,1,1,1,1)$$

$$(3,2,2,1,1,1)$$

$$(2,2,2,2,1,1)$$

After having the all the subspaces according to the unlabelled degree sequence, we can further divide each subspace into smaller subspaces in which all the graphs have the same labelled degree sequence. The number of such subspaces for each unlabelled degree sequence can be calculated by counting the total possible permutations in that unlabelled degree sequence. Assume in an unlabelled degree sequence, there are m_k nodes having degree d_k, therefore $\sum_{k=1}^{K} m_k = n$, where K is the total number of degrees which are different. For example, a star with n nodes has two different degrees: $d_1 = n-1$ and $d_2 = 1$, therefore $m_1 = 1$ and $m_2 = n-1$. The total number of possible labelled degree sequence for the star is $\binom{n}{1} = n$, that is, any node can be the highest degree node. For a general unlabelled degree sequence when the number of nodes having degree d_k is m_k, the total number of labelled degree sequence is

$$\binom{n}{m_1, m_2, \ldots, m_K} = \frac{n}{m_1! m_2! \ldots m_k!}. \tag{5.6}$$

Equation (5.6) shows that if an unlabelled degree sequence has more distinct degrees, it contains more labelled degree sequences. Without tree constraint, the unlabelled degree sequence with n nodes which has the maximum number of labelled degree sequence should be $(n-1, n-2, n-3, \ldots, 2, 1, k)$ where k would be any number between 1 and $n-1$. There is only one repeated degree and the total number of labelled degree sequence is $n!/2$. For trees, although the degree sequence has to be constrained as $\sum d_i = 2(n-1)$, the degree

sequence with maximum number of labelled degree sequence can be estimated roughly as consecutive numbers from 1 to k with repeated degree within $[1, k]$, where k can be estimated from the equation $k(k+1) + (n-k) = 2(n-1)$.

For each labelled degree sequence, van Lint and Wilson [117] calculate the number of labelled tree using the definition of multinomial coefficient and induction. For a degree sequence $d_1, d_2, \ldots, d_n$, the number of labelled tree is:

$$\binom{n-2}{d_1-1, \ldots, d_n-1} = \frac{n-2}{(d_1-1)! \ldots (d_n-1)!}. \tag{5.7}$$

Equation (5.7) indicates that for a tree with n nodes, the maximum number of labelled trees is obtained when the degree sequence satisfies $(2, 2, \ldots, 2, 1, 1)$, while the minimum number of labelled trees is derived at the case when degree sequence is $(n-1, 1, 1, \ldots, 1)$. That is, the chain structure has the maximum number of labelled graphs ($(n-2)!$) while the star structure has the minimum number of labelled graphs (only 1 labelled star for a labelled degree sequence).

The general flip defines the stationary probability of each labelled graph while doing random walks on G1, while the number of labelled graph for a certain degree sequence can be exactly calculated in the previous section. Combining both results, we can derive the relationship between stationary probability of graphs given each degree variability. Although each higher variability labelled graph is more likely to appear than each lower variability graph, it does not necessarily mean that lower variability graphs are less likely because the likelihood also depends on the total number of graphs with that variability. In fact, we can exactly calculate the probability of a graph given an unlabelled degree sequence $(d_1, d_2, \ldots, d_n)$ as:

$$c\left(\sum_{i=1}^{n} d_i - 2l\right) \frac{n-1}{(d_1-1)! \ldots (d_n-1)!} \frac{n}{m_1! m_2! \ldots m_k!} \tag{5.8}$$

where c is a constant to normalize the whole item to be a probability. The second item is proportional to the probability of each labelled tree while doing random walk, the third item is the the number of labelled tree for a given labelled degree sequence, while the last

term is the number of labelled degree sequence given an unlabelled degree sequence. For a given degree sequence, we can calculate the variability of the degree sequence, therefore build relationship between probability of the graphs having that degree variability.

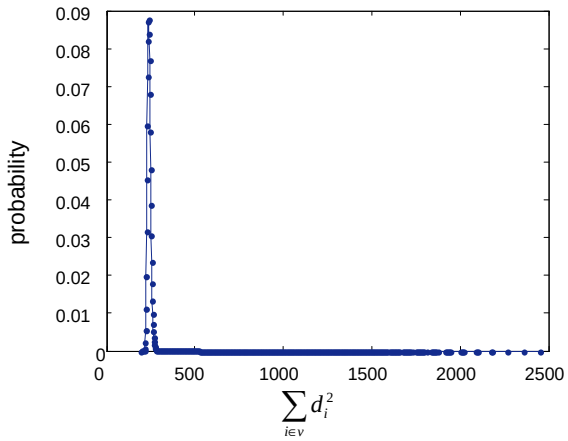

Figure 5.5: Graph probability vs. degree sequence variability.

Figure 5.5 shows probability of graphs having different variability of graph degree sequence for a tree with $n = 50$ nodes (here we use $\sum_{i=1}^{n} d_i$ for degree sequence variability). For each $\sum_{i=1}^{n} d_i$, we can obtain all the degree sequences having that value, and calculate the probability of graphs of each degree sequence and add them up according to different degree variability. In fact, figure 5.5 suggests that the low degree variability graphs are much more likely to appear than the high degree variability ones. The reason is that the third and last terms in equation (5.8) completely dominate the evaluation and these two terms favor low variability graphs much more than high variability ones. For a tree of $n = 50$ nodes, using chain structure as an example of low variability graph, the third term for chain is about $(n-2)! = 1.2 \times 10^{61}$ and the last term is $n(n-1)/2 = 1225$, therefore the total number of labelled chains for 50 nodes is about 1.5×10^{64}. However, for the highest degree variability graph, the star structure, the third term is only 1 and the last term is n. The comparison results are striking, the total number of labelled chain is 3^{62} more that the total number of labelled star for the graphs with same number of nodes and links, yet each labelled star has stationary probability only 20 times bigger than each labelled chain. When counting all the degree sequence with different variability, the lower degree variability graphs are much

more than the higher one, therefore when taking random walk on G1, it is more likely to arrive at lower degree variability graphs than higher ones. Traditional random graph theory focuses on low variability graphs since they are more likely to be generated randomly. This also confirms that the high degree variability graphs in most complex networks exist for special purpose, not from generic and random mechanism. For example, the Internet comes from the highly engineered design.

5.4 The Degree Preserving Flip and G2

The general flip can be regarded as the most fundamental transformation method for simple and connected graphs with the same numbers and links. Any other transformation method in this space or in the subspace can be performed by one or several steps of general flips, since the general flip can change any graph to the other within finite steps in this graphs space. Here we introduce a two-step of general flips, the degree-preserving flip, which can transform graphs in a subspace of G1 where all the graphs have the same labelled degree sequence. We call this subspace as G2, and without surprising the properties of G2 are most relevant to the structural metric, the *s*-metric, introduced before. This reenforce the importance of the *s*-metric to study the graphs having the same degree sequence.

5.4.1 The Degree Preserving Flip

The degree-preserving flip is defined in [79], and it has the nice property that the transformation keeps the same nodal degree. The transformation is performed as follows: consider four nodes $u, v, w, x \in \mathcal{V}$, if they are connected as $(u, v), (v, w), (w, x) \in \mathcal{E}$, and if neither u, w nor v, x is connected, flip links (u, v) and (w, x) to (u, w) and (v, x) (see figure 5.6). If any one of u, w or v, x is connected, the graph remains the same. The domain space under this transformation contains all the simple and connected graphs with the same degree sequence, and furthermore the degree of each node is fixed.

A well-known transformation to preserve the degree distribution of the graph is the degree preserving rewiring, which exchanges any two links (u, v) and (w, x) into (u, w) and

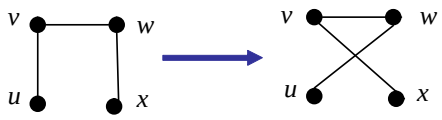

Figure 5.6: The degree-preserving flip.

(v, x) without considering the connectivity between v, w. Our degree-preserving flip, as a special case of the degree preserving rewiring, can maintain the simplicity and connectivity without checking any global information, thus can save a lot of computation.

The soundness and feasibility of the degree-preserving flip for a regular graph where all the nodes have exactly the same degrees have been proven in [79]. This can easily be extended to nonregular graphs with any degree sequence. The generality of the degree-preserving flip for regular graphs is also proved in [79], which however is hard to extend to nonregular graphs. In fact, it has been shown that the degree-preserving flip for nonregular graphs may not be general for some degree sequences and a counterexample named a bow-tie switch is provided in [50]. In a bow-tie graph, nodes x, y, z, t, v form a link set (x, y), (z, t), (x, v), (y, v), (z, v), (t, v), where v is the center of the bow-tie graph and the other nodes are edge nodes. The space of all graphs with this degree sequence contains only two graphs (see figure 5.7), however, these two graphs cannot reach each other by the degree-preserving flip. The transformation between these two bow-tie graphs is called the bow-tie switch.

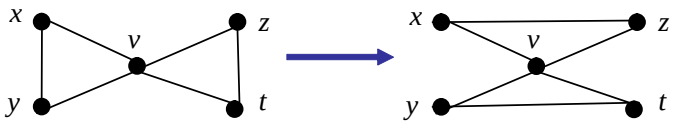

Figure 5.7: The bow-tie switch.

Despite of the existence of bow-tie graphs, we prove that the degree-preserving flip is general by adding a small constraint to the graph. In fact, it has been proved in [50] that the degree-preserving flip is general when the graph has a diameter greater than 3. We can further relax this constraint.

Lemma 5.4.1. *The degree-preserving flip is general when any graph in the domain space has a diameter greater than 2.*

Feder et al. [50] have shown that if adding bow-tie switch, the degree-preserving flip can reach any graph. Based on this result, we show that when the diameter of any graph is greater than or equal to 3, any bow-tie switch can be resolved by a sequence of degree-preserving flips therefore any graph is reachable.

Proof. Assume that a graph contains a subgraph which is a bow-tie graph. Since a bow-tie structure has a diameter equal to 2, there must exist nodes and links other than those in the bow-tie subgraph to make the graph diameter greater than or equal to 3. The extra nodes may connect to any of the edge nodes x, y, z, t, or the center node v, resulting in two basic scenarios. Case 1 happens when at least one node other than these in bow-tie graph is connected to one or mode edge nodes, yet is not connected to the center node. We assume that a node u is connected to any or both of x, y, but is not connected to any of z, t, v (otherwise the diameter of the graph may not great than 2). We can perform three degree-preserving flips to resolve the bow-tie switch as shown in figure 5.8.

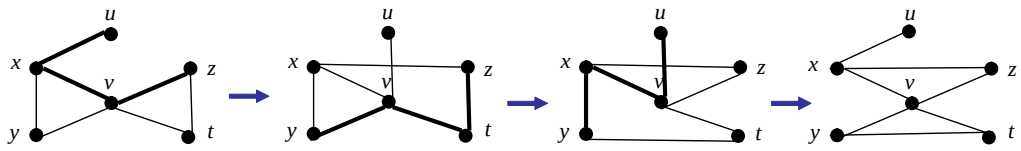

Figure 5.8: The degree-preserving flip for bow-tie switch when extra node is connected to the edge nodes of bow-tie (case 1).

In the second case when there is no node connected to the edge nodes and all the extra nodes are connected to the center node v. To ensure the diameter of the graph greater than or equal to 3, there must exist two extra nodes u, w connected as a chain and one of which is connected to the center node v such that $(u, v), (u, w) \in \mathcal{E}$ and w does not connect to v. Still, we can perform six steps of degree-preserving flips to get the bow-tie switch as shown in figure 5.9.

For any graph which contains a bow-tie with a diameter greater than or equal to three, we claim these graphs can be transformed to either of the two basic cases. For example, in

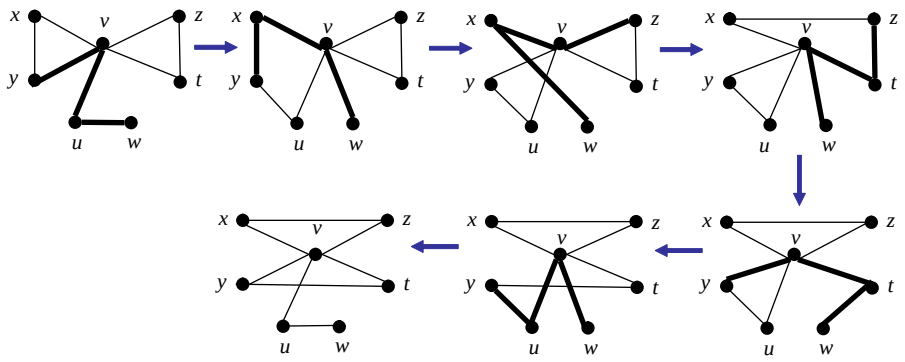

Figure 5.9: The degree-preserving flip for bow-tie switch when extra nodes are connected to the center node (case 2).

figure 5.4, node u can also be connected to node y and this will not affect the flip process at all. If u is also connected to z, we need another node connecting to u to make the diameter greater than two. In this case we can flip the extra link and link (u, z) to release the connection between u and z so as to unblock the triangle for a new flip. □

When the graph has a diameter equal to two, no degree-preserving flip can happen. When the graph has a diameter great than two, the degree-preserving flip can not reduce the diameter to two, since the degree-preserving flip is reversible. In fact, when the diameter is greater than two, by definition there must exist two nodes between which the shortest path is greater than two. Along this shortest path, there is no extra link connecting any node on the path, therefore doing the degree-preserving flip on this path cannot be blocked. The removal of the existing links will free any triangle that is formed by these links, and a new degree-preserving flip can now be performed.

5.4.2 Properties of G2

G2 is defined by the degree-preserving flip, such that each node represents a simple and connected graph with the same degree sequence as the other graphs in G2, and two nodes share a same link if and only if they can be transformed by one degree-preserving flip. It turns out that G2 has some nice properties related to the s-metric. Since the s-metric is a measure of the extend to which a graph is scale-free, scale-free graphs are special in G2.

Lemma 5.4.2. *The degree of each node in G2 is* $s - \sum_{i=1}^{n} d_i^2 + l - \delta$*, where* δ *is the number of triangles which will be specially defined later.*

The degree of each node in G2 depends on how many degree-preserving flips that can happen for the corresponding graph. This is determined by the number of 3-motifs in the graph and the number of triangles consisting two links which belongs to the 3-motifs. A 3-motif is a subgraph that contains four nodes u, v, w, x which are connected as $(u, v), (v, w), (w, x)$. The total number of 3-motifs is (see Section 4.2 for details)

$$\sum_{(i,j)\in\mathcal{E}} (d_i - 1)(d_j - 1) = \sum_{(i,j)\in\mathcal{E}} d_i d_j - \sum_{i=1}^{n} d_i^2 + l = s - \sum_{i=1}^{n} d_i^2 + l.$$

When u, w or v, x are connected, the degree-preserving flip cannot happen. This case can be counted as the triangle in the 3-motif. Denote the total number of 3-motifs that contain triangle as δ. Notice this δ is different from the previous Δ which is defined as the total number of triangles in the graph since a 3-motif may contain 2 triangles which should be counted as one for this motif. From the total number of 3-motifs and δ, we can derive that the degree of each node in G2 is

$$s - \sum_{i=1}^{n} d_i^2 + l - \delta. \tag{5.9}$$

Since all the graphs in G2 have fixed degree sequences, the degree of node in G2 is determined by the s-value, therefore scale-free graphs have more neighbors than the other graphs.

Lemma 5.4.3. *When performing a random walk from any graph on G2 with equal probability to all its neighbors, the final stationary distribution for staying at this graph is proportional to* $s - \sum_{i=1}^{n} d_i^2 + l - \delta$*.*

Similar argument as before, the stationary distribution of a node in a graph is proportional to its degree. An important result from this lemma is that the scale-free graphs are more likely to appear in the space of graphs with the same degree distribution when performing a random walk in the space of graphs.

5.5 Summary

The space of graphs is fully explored in this section by introducing graph transformation methods among graphs, particularly those with the same numbers of nodes and links, and those with the same degree sequence. The general flip, a most fundamental transformation method, defines G1 which contains all the simple and connected graphs with the same numbers of nodes and links. G2 is a subspace of G1 containing the graphs with the same degree sequence, and the graph relationship is built from the degree-preserving flip, a two steps general flip. We characterize the properties of both G1 and G2, and prove that when fixing the numbers of nodes and links, G1 has a nice relationship to the variability of the degree sequence of each graph in G1, when further constrained to the graphs with the same degree sequence, many properties of G2 are related to the *s*-metric. By exploring the space of graphs, we obtain a much clearer picture on the number of graphs having different degree sequences, and on the fundamental differences and similarities among those graphs.

Chapter 6

Other Projects

My research spans both horizontally and vertically. The majority of this thesis elaborates my horizontal interests: the topologies of complex networks, where we consider the role of functions and constraints for the Internet router-level topology, define a structural metric to differentiate graph models with the same degree sequence in general complex networks, and study the space of graphs for a clearer picture of graph relationship and transformations. My vertical line builds on the top of the Internet topology and extends to the protocol stack of the Internet, including the investigation of joint optimization of routing and transferring rate at the IP and TCP layers [119], a theoretical framework for Internet congestion control for TCP layer [101], and user-perceived failure detection using packet trace at application layer[72]. In this chapter, these three projects along the vertical line will be briefly described.

6.1 Cross-Layer Optimization in TCP/IP Networks

Recent studies have shown that any TCP congestion control algorithm can be interpreted as carrying out a distributed primal-dual algorithm over the Internet to maximize aggregate utility, and a user's utility function is defined by its TCP algorithm, see e.g. [66, 74, 88, 83, 76, 68, 73] for unicast, [64, 38] for multi-cast, and [75, 65, 114] for recent surveys and further references. All of these works assume that routing is given and fixed at the timescale of interest, and TCP, together with active queue management (AQM), attempt to maximize aggregate utility over source rates. In this chapter, we study the cross-layer

utility maximization at the timescale of route changes.

We focus on the situation where a single minimum-cost route (shortest path) is selected for each source-destination pair. This models IP routing in the current Internet within an Autonomous System using common routing protocols such as OSPF [91][1] or RIP [58]. Routing is typically updated at a much slower timescale than TCP–AQM. We model this by assuming that TCP and AQM converge instantly to equilibrium after each route update to produce source rates and "congestion prices" for that update period. These congestion prices may represent delays or loss probabilities across network links. They determine the next routing update in the case of dynamic routing, similar to the system analyzed in [53]. Thus TCP–AQM/IP form a feedback system where routing interacts with congestion control in an iterative process. We are interested in the equilibrium and stability properties of this iterative process. To simplify notation, we will henceforth use TCP–AQM/IP and TCP/IP interchangeably.

Here are our main results. In the case of pure dynamic routing, i.e., when link costs are the congestion prices generated by TCP–AQM, it turns out that we can interpret TCP/IP as a distributed primal-dual algorithm to maximize aggregate utility over *both* source rates (by TCP–AQM) and routes (by IP) if it TCP/IP converges. We consider the problem, and its Lagrangian dual, of maximizing utility over source rates and over routing that use only a *single* path for each source-destination pair. Unlike the TCP-AQM problem or the multi-path routing problem that are convex optimizations with no duality gap, the single path TCP/IP problem is non-convex and generally has a duality gap. Equilibrium of the TCP/IP system exists if and only if this problem has no duality gap. In this case, TCP/IP equilibrium solves both the primal and the dual problem. Moreover, it incurs no penalty for not splitting traffic across multiple paths: optimal single-path routing achieves the same aggregate utility as optimal multi-path routing. Multi-path routing can achieve a strictly higher utility only when there is a duality gap between the single-path primal and dual problems, but in this case, the TCP/IP iteration does not even have an equilibrium, let alone solving the utility maximization problem.

[1]Even though OSPF implements a shortest-path algorithm, it allows multiple equal-cost paths to be utilized. Our model ignores this feature.

Even when the single-path problem has no duality gap and TCP/IP has an equilibrium, the equilibrium is generally unstable under pure dynamic routing. It can be stabilized by adding a sufficiently large static component to the definition of link cost. The existence and characterization of TCP/IP equilibrium when the link costs are not pure congestion prices, however, are open problems. To proceed, we specialize to a ring network with a common destination and demonstrate an inevitable tradeoff between utility maximization and routing stability. Specifically, we show that the TCP/IP system over the special ring network is indeed unstable when link costs are pure prices. It can be stabilized by adding a static component to the link cost, but at the expense of a reduced utility in equilibrium. The loss in utility increases with the weight on the static component. Hence, while stability requires a small weight on prices, utility maximization favors a large weight. We present numerical results to validate these qualitative conclusions in a general network topology. They also suggest that routing instability can reduce aggregate utility to less than that achievable by (the necessarily stable) pure static routing.

Indeed we show that if the link capacities are optimally provisioned, then *pure static* routing is enough to maximize utility even for general networks. Moreover, it is optimal within the class of multi-path routing: again, there is no penalty at optimality in not splitting traffic across multiple paths.

The duality model of TCP–AQM has been useful in understanding the equilibrium properties, including throughput, packet loss, delay, and fairness, of large-scale networks under TCP–AQM control. This work is a first, and preliminary, attempt to apply the same methodology to understand the cross-layer interaction of TCP–AQM, minimum-cost routing and resources allocation. Our model is simplistic – it ignores finite duration flows and randomness in real networks, and reduces the rich behavior of IP to minimum-cost routing. Even within this highly abstract model, many questions remain open. First, even though numerical examples suggest that the tradeoff between routing stability and utility maximization is present in a more general network than the special ring network we studied, we have not been able to find an analytical proof. One of the major difficulties is that, in a general network, minimum-cost routing cannot be as conveniently represented as in the ring network. Second, when static component is included in link cost, it is not known if

TCP/IP has an equilibrium, whether the equilibrium jointly solves a certain optimization problem, and under what condition it is stable. Third, it would be interesting to estimate the duality gap in the single-path problem. Even though this problem is not directly related to the TCP/IP iteration when the duality gap is nonzero, the gap measures the penalty of not splitting traffic among multiple paths.

6.2 Methodological Frameworks For Internet Congestion Control

In this work, we make the theme that optimization based decompositions of complex systems into interacting modules facilitates analysis, comparability and verifiability of the desired system properties. The modularity that such decompositions offer, and which at first endows the systems with an apparent complexity should be taken advantage of. Aiming for such decompositions is beneficial both for analysis and design.

We also stress that in any analysis procedure it is important to construct robust models for the modules, as this will capture the uncertainty in modelling and component parameters so that it be taken account in the design process. The new tools that we develop in this work allow us to analyze such systems even at the nonlinear level, and expand the applicability of this methodology.

Complex systems and large scale networks will dominate the future societies as technology advances. Designing such systems is more than art based on intuition. It is widely appreciated that network congestion control for the Internet is probably the only complex system for which we have a good understanding of the interaction of the various modules at the TCP/AQM level. The system can be designed by resorting to a solid methodological framework that provides the desired functionality at equilibrium, based on an optimization scheme; and the correct dynamics can be chosen for the various modules to drive the system to the equilibrium – the right choice of dynamics are key to the scalability of the verification result.

The success in designing network congestion control schemes for the Internet through

a mathematical formulation which enables understanding of its functionality and the limitations that features such as delays pose, allows us to believe that similar hierarchical structures can enable understanding and design of other complex systems in the future. We envision that solid methodological frameworks can be used to formulate and solve the design problem this way and the resulting system's functionality can be proven in a structured way. Apart from the specific analysis results that one can produce by hand, the algorithmic procedure we propose can be used to analyze more complicated system descriptions therefore increasing the set of model building blocks that can be used in the construction of future mathematical frameworks for complex system analysis.

6.3 Detect User Perceived Faults Using Packet Traces

Fault detection in a timely fashion is critical for network management. In this work, we concentrate on a specific variant - given a packet trace from any link in an edge network (such as a university network or an enterprise network), detect faults that can be perceived by the end user. Identifying significant faults is an essential first-step to localize and fix the faults, a problem that is often complicated by the fact that human users rarely report faults and even when they do, human reports tend to be vague and unreliable. Our goal is to build a tool that processes packet traces online, identifies application level faults that will be perceived by an end-user (marked slowdown, incorrect response, disconnection etc.) and raises an alarm to the fault localization system so that failure can be detected without human interference. Though clearly desirable, the wide variety of applications, protocols and unavoidable low-level nitty-gritty of real edge networks makes it challenging to build such a tool.

A general definition of failure is the condition that "the delivered service deviates from the specified service [69]." There is much prior work attempts to detect faults, but either at a too-coarse granularity where all traffic is on an Internet path or at a too-fine granularity which considers the performance of a single TCP flows. Our focus is different, we want to identify faults that matter to real users. In edge networks, each user transaction may involve many flows, connections to many different servers that may traverse different paths.

We propose a mechanism to detect the failure perceived by the end user from packet traces collected at the end user side. It is a passive approach in the sense that it does not add any additional packets to the network just for the purpose of detection. While active probing is considered as a more accurate approach, it introduces extra packets into the network and overloads the servers and networks. Furthermore different protocol requires different active probing method therefore it is hard to generalize to other application protocols. Passive monitoring can reduce the network overload and it is easy to be implemented in the server side, user side or in the network. Since it is able to collect all the packets from a user to the outside network, it essentially can be used to reconstruct the protocols if we are able to parse into the packet information, therefore it could be generalized to detect fault for different application protocols.

Detection failure from packet traces has proved to be a challenging problem. Since the packet trace is a collection of all the packets sent and received by the end user, it contains too much information which does not lend itself naturally for inspection. For example, a simple click of a website may contain hundreds of packets which involve several different application protocols such as DNS, WINS, IPSEC, Kerberos, HTTP as well as several different servers (DNS server, WINS server, authentication server, Web server). The development of abstract yet informed models from packet traces to leverage failure information from standard protocol information is crucial. However abstraction of a standard application protocol is not obvious since there are numerous different application protocols (such as HTTP, SMTP, SMB, DNS) which may behave quite differently, also there are many application programs using the same application protocol (for example: internet explorer, Firefox, Mozilla are all based on HTTP protocols), furthermore even for the same application programs, different configuration parameters may result in quite different pattern of packets.

Another challenge of fault detection for end user is that it lacks for a good validation tool. Ideally, we want to compare the results from our detection scheme with a report from end users on every failure they encounter. However obtaining such report is not an easy task since currently there are no such network programs that can automatically generate such report. Requiring end users to manually input the failures takes a lot of work and the

report itself may not be reliable.

Despite of these difficulties, we make consistent progress for fault detection in edge networks. We define failures which accurately reflect actual performance deviation or degradation experienced by end users, and create a robot that can mimic end user behavior of fetching some web-sites periodically and obtain the return information. The failure of the return information is in fact the information that an end user experiences when he/she fetches the same web-site at the same time. Our robot can provide a reliable validation tool for HTTP algorithm.

We utilize protocol dependence and group packets according to their 5-tuple flow information (source IP, destination IP, source port, destination port and protocol). We implement a white-box scheme to detect HTTP failure where we are able to parse into HTTP header information such as command (get, post), HTTP return code (200 OK, or 502 bad gateway). Compared with the robot result, the white-box approach shows very low false positive and false negative rate. Furthermore, robot can only fetch the website that we tell it to do, yet our white-box approach essentially can report the failure happening to all the web-sites that users are browsing, therefore we can remove robot probing and use white-box scheme as a validation tool for other HTTP detection schemes.

We also generalize the white-box approach for HTTP protocols by relaxing the requirement of looking at HTTP header information so that it can deal with the encrypted packets. We aggregate packets according to a user task and use the number of packets, bytes and transaction time to detect failure. Small number of bytes is an indication of content error that HTTP returns such as 401 unauthorized, or 500 internal server error, while the number of packets together with transaction time will reflect the network performance degradation such as link congestion.

Furthermore, we consider a completely different protocol, RPC protocol used by email which requires consistent exchanging of information between users and servers. We characterize fault by a large number of small packets and we are able to detect the OUTLOOK failure such as unable to connect to the server. The fault detections described herein are not limited to any particular application layer or transport layer, but may be applied to many contexts and environments.

Finally compared with the failure detected from commercial fault detection products for edge networks such as Microsoft Operation Management, our failure detection scheme can detect user perceived faults in an enterprise network in a much more reliable and accurate manner. Traditional fault detection tools do not consider end users, therefore are inadequate for localizing performance faults, such as identifying users are dissatisfied with the end-to-end response time.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

The search for unifying properties of complex networks is popular, challenging, and important. It is certainly appealing that scale-free network models can avoid all the domain specific details for different complex networks, yet make interesting and testable predictions. Unfortunately, this fact yields results that collapse when tested with elaborated analysis on the functionality of the Internet, as well as structural analysis of graphs having the same high variability degree distributions.

In this thesis, using Internet as a case study, we have shown that there exist technological, economic, and graph theoretic reasons why the most popular scale-free models cannot be true when they are used to describe current Internet router-level topology. We propose a complementary approach of combining a more subtle use of statistics and graph theory with a first-principles theory of router-level topology that reflects practical constraints and trade-offs. While there is an inevitable trade-off between model complexity and fidelity, a challenge is to distill from the seemingly endless list of potentially relevant technological and economic issues the features that are most essential to a solid understanding of the intrinsic fundamentals of network topology. We can successfully address this challenge by providing a Heuristic Optimal Topology (HOT) model that incorporates hard technological constraints on router bandwidth and link connectivity, together with abstract models of user demand and network performance. In a high performance and low cost network, the high bandwidth core router cannot have high degree due to the router technological constraint,

while end-user demands and economic constraints on link costs push all the high variabilities in edge routers. In contrast, scale-free models provide a relatively easy way to generate the desired power-law degree distribution, however their highly connected hubs have such bad performance as to make it completely unrealistic that they could reasonably represent a highly engineered system like an ISP network or the Internet as a whole.

The space of graphs to describe complex network is extremely diverse even within the graphs having the same degree distributions. Network performance provides an important metric to measure the functionality of the Internet. Yet, a subtle look at the structures of scale-free models and the HOT model reveals a fundamental difference between these models and this difference can be captured by a structural metric, the *s*-metric which we introduce to differentiate between all simple, connected graphs having an identical degree sequence, especially when that sequence satisfies a power-law relationship. Elaborating the features of graphs with high or low *s*-metric provides enhanced understanding towards a theory of scale-free networks. The *s*-metric, as a measure of the extend to which a graph has highly connected hubs, together with power-law degree distribution, provides a quantitative definition of scale-free networks. We provide evidence that high *s*-value graphs, i.e., scale-free graphs actually share a wide range of emergence features, such as hublike cores, high likelihood under variety of random generation mechanisms, and various kinds of self-similarity. We also suggest that when making statements about a graph based on these properties one must consider the background set against which these properties are being evaluated.

While the functional and structural metrics provide two effective views of looking at the highly dimensional space of graphs, we furthermore introduce a new paradigm to understand the space of graphs by building connections between two isolated graphs according to some local transformations. Exploring this connected space of graphs gives a cleared picture of this space and reenforces the important role that the variability of degree sequence plays in the graphs with the same numbers of nodes and links, and the role of the *s*-metric in the subspace of graphs having the same degree sequence.

7.2 Future Directions

There are several future works that can directly outgrow from this thesis.

7.2.1 Internet Topology Generator

"All models are wrong, but some are useful." — G. P. E. Box.

Any work on Internet topology generation and evaluation runs the danger of being viewed as incomplete and/or too preliminary if it does not deliver the "ultimate" product, a topology generator. A natural extension of our work is to build a useful topology generator so that protocols can be evaluated before being implemented to the real network. Our HOT model opens up a new line of the Internet research in identifying causal forces that are either currently at work in shaping large-scale network properties or could play a critical role in determining the lay-out of future networks. Our model is still in a toy model stage, so called because it only leverages the most important aspects of the Internet router-level topology, and therefore provides the most coarse-grained level topology generator. More functional requirements and physical constraints can be added to our approach, which results in a higher model fidelity but at a cost of higher model complexity.

An ideal topology generator should allow one to incorporate different level of details and produce different topologies that address these details. For example, when taking network reliability into consideration, many gateway routers or edge routers should be multihomed. That is, instead of the tree structure from core to edge routers, one edge router should be connected to two or more higher level routers, to ensure that when one link or router is broken, it can still connect to the whole network. The second example is that in the case when the traffic matrix is given and not consistent with our gravity model, we should adjust the corresponding connections from the edge routers which directly carry the traffic to higher level routers so as to maximally utilize the router capacity. A more complete example should also consider more possible constraints such as router geographical location, router cost when operating at different technological boundaries, link length cost, political reasons, etc., as well as other functional requirements, such as low latency, link or router redundancy. While the optimization problem that incorporates all these constrains and ob-

jectives is almost impossible to solve theoretically, we can take an approach based on our HOT toy model and do local adjustment or optimization to improve or fulfill the function requirement within the constraints. Despite all these different variations, we expect that any Internet router-level topology generated according to constraints and objectives should have sharp difference from scale-free networks, no matter characterized by performance metric or the *s*-metric.

It is widely recognized that performance of the same protocol can be quite different under different topologies that run on top of it. For example, many TCP protocols are stable in the single bottleneck link case while oscillate in the multiple bottleneck links case. However, how different topologies would affect our protocol evaluation is not entirely understood. A theoretical analysis could be quite complicated if not impossible. With the aids of the topology generator, we can carry out this analysis in an empirical way. That is, protocols can be evaluated under different topologies, or similar topologies with subtle transformations. A topology generator which performs many different protocol evaluations can in turn suggest improvements in network design, leading to well protected, reliable and high performance next generation networks.

7.2.2 Apply the HOT Idea to Other networks

We do not claim that the results obtained for the router-level topology of the Internet pertain to other complex networks. However, even for these completely different cases, we believe that methodologies that explicitly account for relevant functionality and constraints, or other key aspects can provide similar insight into what matters when understanding, or evaluating the corresponding topologies. A detail understanding of the functions and constraints of the network is important, since it is exactly these functions and constraints which drive the underlying structure of networks. For example, we would expect that the AS (autonomous system, which represents an administrative domain such as a company or a school) level topology of the Internet would be quite different from the router-level topology, since at AS-level, the major concern is that how each domain sends it traffic to maximize its own profit while minimize its own cost. Here the connection between

each domain is more like business relationship. In this problem, game theory, or more specifically, industrial organization will provide a necessary tool to obtain the heuristic optimal topology.

7.2.3 Benchmark Graphic Metrics

A rough estimation on the number of existing metrics to evaluate graphs easily yields an order of a hundred. Although in our thesis we claim that the functional metric is an essential consideration when designing the Internet router-level topology, and the structural metric plays an important role to differentiate graphs with the same degree distribution, a systematic way to evaluate the effectiveness graphic metrics requires much more work.

On one side, same metric which is used to measure the properties of one network many not be effective for other networks due to the complications and highly diversity of graphs. Our investigation of assortativity shows that some metrics would be extremely misleading without considering the background they apply to. A good metric should give the same qualitative answer, invariant in different background set, or explicitly state the related background set it compares against when it is meaningful. In fact, we conjecture that many properties which scale-free networks have in common with real networks are resulted from the high variability of degrees instead of hublike core, yet the former does not necessarily implies the latter. For example, the average shortest path lengths of HOT model and scale-free model are both nicely low, due to the aggregation at the high degree nodes. However concluding that the hublike core scale-free model is representative from this evidence would be a mistake.

On the other side, many graph metrics may have some common properties except for different scenarios (for example, the s-metric and assortativity). Categorizing all these metrics according to their fundamental similarities and differences will greatly clarify the current literature and facilitate future studies.

7.2.4 the Graph Space: Properties and Dynamics

While the space of graphs is extremely diverse and the connected space of graphs is much more complicated than the graph it contains, our work sheds light on this highly dimensional space by dividing it into countable subspaces and by relating its node degree, stationary distribution to the graph properties that each node represents. It would be of great interests to further explore each subspace of graphs, so that more fundamental similarity and difference between each graphs can be easily characterized. We can also capture the change of functionality along a series of structural transformations towards a better understanding of the interactions between function and structure.

Another important application for this GRAPH of graph is to characterize the dynamics of graphs. As flips define a local and gradually changing process for graph transformation, we can evaluate how functionalities and structures changes along this process, especially when the transformations are along the directions of monotonically changing the degree variability, or the s-metric. Furthermore, we can guide dynamics according to certain functional requirements and practical constraints to study the evolution of complex networks.

Appendix A

Constructing an $s_{\max}$ Graph

As defined previously, the $s_{\max}$ graph is the element g in some background set G whose connectivity maximizes the quantity $s(g) = \sum_{(i,j)\in\mathcal{E}} d_i d_j$, where d_i is the degree of node $i \in \mathcal{V}$, $\mathcal{E}$ is the set of links that define g, and $D = \{d_1, d_2, \ldots, d_n\}$ is the corresponding degree sequence. Recall that since D is ordered according to $d_1 \geq d_2 \geq \cdots \geq d_n$, there will usually be many different graphs with nodes satisfying D. The purpose of this appendix is to describe how to construct such an element for different background sets, as well as to discuss the importance of choosing the "right" background set.

A.1 Among "Unconstrained" Graphs

As a first case, consider the set of graphs having degree sequence D, with only the requirement that $\sum_{i=1}^{n} d_i$ be *even*. That is, we do not require that these graphs be simple (i.e., they can have self-loops or multiple links between nodes) or that they even be connected, and we accordingly call this set of graphs "unconstrained." Constructing the $s_{\max}$ element among these graphs can be achieved trivially, by applying the following two-phase process. First, for each node i: if d_i is even, then attach $d_i/2$ self-loops; if d_i is odd, then attach $(d_i - 1)/2$ self-loops, leaving one available "stub." Second, for all remaining nodes with "stubs," connect them in pairs according to decreasing values of d_i. Obviously, the resulting graph is not unique as the $s_{\max}$ element (indeed, two nodes with the same degree could replace their self-loops with connections among one another). Nonetheless, this construction does maximize $s(g)$, and in the case when d_i is even for all $i \in \mathcal{V}$, one achieves an $s_{\max}$ graph with

$s(g) = \sum_{i=1}^{n}(d_i/2)\cdot d_i^2$. As discussed in section 4.4, against this background of unconstrained graphs, the $s_{\max}$ graph is the perfectly assortative (e.g., $r(g) = 1$) graph. In the case when some d_i are odd, then the $s_{\max}$ graph will have a value of $s(g)$ that is somewhat less and will depend on the specific degree sequence. Thus, the value $\sum_{i=1}^{n}(d_i/2) \cdot d_i^2$ represents an idealized upper bound for the value of $s_{\max}$ among unconstrained graphs, but it can only be realized in the case when all nodal degrees are even.

A.2 Among Graphs in $G(D)$

A significantly more complicated situation arises when constructing elements of the space $G(D)$, that is, simple connected graphs having n nodes and a particular degree sequence D. Even so, not all sequences D will allow for the connection of n nodes, i.e., the set $G(D)$ may be empty. In the language of discrete mathematics, one says that a sequence of integers $\{d_1, d_2, \ldots, d_n\}$ is *graphical* if it satisfies the degree sequence of some simple, connected graph, that is if $G(D)$ is nonempty. One characterization of whether or not a sequence D corresponds to a simple, connected graph is due to Erdös and Gallai [46] as discussed in 4.1.

Our approach to constructing the $s_{\max}$ element of $G(D)$ is via a heuristic procedure that incrementally builds the network in a greedy fashion, by iterating through the set of all potential links $\mathcal{O} = \{(i, j) : i < j; i, j = 1, 2, \ldots, n\}$, which we order according to decreasing values of $d_i d_j$. In what follows we refer to the value $d_i d_j$ as the *weight* of link (i, j). We add links from the ordered list of elements in $\mathcal{O}$ until all nodes have been added and the corresponding links satisfy the degree sequence D. To facilitate the exposition of this construction, we introduce the following notation. Let $\mathcal{A}$ be the set of nodes that have been added to the partial graph $\tilde{g}_{\mathcal{A}}$, such that $\mathcal{B} = \mathcal{V}\backslash\mathcal{A}$ is the set of remaining nodes to be added. At each stage of the construction, we keep track of the *current degree* for node i, denoted $\tilde{d}_i$, so that it may be compared with its *intended degree* d_i (note that $\tilde{d}_i = 0$ for all $i \in \mathcal{B}$). Define $\tilde{w}_i = d_i - \tilde{d}_i$ as the number of remaining *stubs*, that is, the number of connections still to be made to node i. Note that values of $\tilde{d}_i$ and $\tilde{w}_i$ will change during the construction process, while the intended degree d_i remains fixed. For any point during

the construction, define $\tilde{w}_{\mathcal{A}} = \sum_{i \in \mathcal{A}} \tilde{w}_i$ to be the total number of remaining stubs in $\mathcal{A}$ and $d_{\mathcal{B}} = \sum_{i \in \mathcal{B}} d_i$ to be the total degree of the unattached nodes in $\mathcal{B}$. The values $\tilde{w}_{\mathcal{A}}$ and $d_{\mathcal{B}}$ are critical to ensuring that the final graph is connected and has the intended degree sequence. In particular, our algorithm will make use of several conditions.

Condition A-1: (Disconnected Cluster). If at any point during the incremental construction the partial graph $\tilde{g}_{\mathcal{A}}$ has $\tilde{w}_{\mathcal{A}} = 0$ while $|\mathcal{B}| > 0$, then the final graph will be disconnected.

Proof. By definition $\tilde{w}_{\mathcal{A}}$ is the number of stubs available in the partial graph $\tilde{g}_{\mathcal{A}}$. If there are additional nodes to be added to the graph but no more stubs in the partial graph, then any incremental growth can occur only by forming an additional, separate cluster. □

Condition A.1a: (Disconnected Cluster). If at any point during the construction algorithm the partial graph $\tilde{g}_{\mathcal{A}}$ has $\tilde{w}_{\mathcal{A}} = 2$ with $|\mathcal{B}| > 0$, then adding a link between the two stubs in $\tilde{g}_{\mathcal{A}}$ will result in a disconnected graph.

Proof. Adding a link between the two stubs will yield $\tilde{w}_{\mathcal{A}} = 0$ with $|\mathcal{B}| > 0$, thus resulting in condition A.1. □

Condition A.2: (Tree Condition). If at any point during the construction

$$d_{\mathcal{B}} = 2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}, \tag{A.1}$$

then the addition of all remaining nodes and links to the graph must be *acyclic* (i.e., tree-like, without loops) in order to achieve a single connected graph while satisfying the degree sequence.

Proof. To see this more clearly, suppose that for some intermediate point in the construction process that $\tilde{w}_{\mathcal{A}} = m$. That is, there are exactly m remaining stubs in the connected component to which the remaining nodes in $\mathcal{B}$ must attach. We can prove that, in order to satisfy the degree sequence while maintaining a single connected graph, each of these m stubs must become the root of a tree. First, recall from basic graph theory that an acyclic graph connecting n nodes will have exactly $l = n - 1$ links. Define $\mathcal{B}_j \subset \mathcal{B}$ for $j = 1, \ldots, m$

to be the subset of remaining nodes to be added to stub j, where $\bigcup_{j=1}^{m} \mathcal{B}_j = \mathcal{B}$. Further assume for the moment that $\bigcap_{j=1}^{m} \mathcal{B}_j = \emptyset$, that is, each node in $\mathcal{B}$ connects to a subgraph rooted at one and only one stub. Connecting the nodes in $\mathcal{B}_j$ to a subgraph rooted at stub j will require a minimum of $|\mathcal{B}_j|$ links (i.e. $|\mathcal{B}_j| - 1$ links to form a tree among the $|\mathcal{B}_j|$ nodes plus one additional link to connect the tree to the stub). Thus, in order to connect the nodes in the set $\mathcal{B}_j$ as a tree rooted at stub j, we require $\sum_{k\in\mathcal{B}_j} d_k = 2|\mathcal{B}_j| - 1$, and to attach all nodes in $\mathcal{B}$ to the m stubs we have

$$\begin{aligned} d_{\mathcal{B}} &= \sum_{i\in\mathcal{B}} d_i = \sum_{j=1}^{m} \sum_{k\in\mathcal{B}_j} d_k \\ &= \sum_{j=1}^{m} \left(2|\mathcal{B}_j| - 1\right) \\ &= 2|\mathcal{B}| - m \\ &= 2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}. \end{aligned}$$

Thus, at the point when (A.1) occurs, only trees can be constructed from the remaining nodes in $\mathcal{B}$. □

A.2.1 The Algorithm

Here, we introduce the algorithm for our heuristic construction and then discuss the conditions when this construction is guaranteed to result in the $s_{\max}$ graph.

- Step 0 (Initialization):
  Initialize the construction by adding node 1 to the partial graph; that is, begin with $\mathcal{A} = \{1\}$, $\mathcal{B} = \{2, 3, \ldots, n\}$, and $O = \{(1, 2), \ldots\}$. Thus, $\tilde{w}_{\mathcal{A}} = d_1$ and $d_{\mathcal{B}} = \sum_{i=2}^{n} d_i$.

- Step 1 (Link Selection): Check to see if there are any *admissible* elements in the ordered list O.

  (a) If $|O| = 0$, then Terminate. Return the graph $\tilde{g}_{\mathcal{A}}$.

  (b) If $|O| > 0$, select the element(s), denoted here as (i, j), having the largest weight $d_i d_j$, noting that there may be more than one of them. For each such link (i, j),

check $\tilde{w}_i$ and $\tilde{w}_j$: If either $\tilde{w}_i = 0$ *or* $\tilde{w}_j = 0$ then remove (i, j) from $\mathcal{O}$.

(c) If no admissible links remain, return to STEP 1(a).

(d) Among all remaining links having *both* $\tilde{w}_i > 0$ and $\tilde{w}_j > 0$, select the element (i, j) with the largest value $\tilde{w}_i$ (where for each (i, j) $\tilde{w}_i$ is the *smaller* of $\tilde{w}_i$ and $\tilde{w}_j$), and proceed to STEP 2.

- STEP 2 (LINK ADDITION): For the link (i, j) to be added, consider two types of connections.

  – Type I: $i \in \mathcal{A}$, $j \in \mathcal{B}$. Here, node i is the highest-degree node in $\mathcal{A}$ with non-zero hubs (i.e., $d_i = \max_{k\in\mathcal{A}} d_k$ and $\tilde{w}_i > 0$) and j is the highest-degree node in $\mathcal{B}$. Add link (i, j) to the partial graph $\tilde{g}_{\mathcal{A}}$: remove node j from $\mathcal{B}$ and add it to $\mathcal{A}$, decrement $\tilde{w}_i$ and $\tilde{w}_j$, and update both $\tilde{w}_{\mathcal{A}}$ and $d_{\mathcal{B}}$ accordingly. Remove (i, j) from the ordered list $\mathcal{O}$.

  – Type II: $i \in \mathcal{A}$, $j \in \mathcal{A}$, $i \neq j$. Here, i and j are the largest nodes in $\mathcal{A}$ for which $\tilde{w}_i > 0$ and $\tilde{w}_j > 0$.

    * Check the *Tree Condition*:
      If $d_{\mathcal{B}} = 2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}$, then Type II links are not permitted. Remove the link (i, j) from $\mathcal{O}$ *without adding it to the partial graph*.
    * Check the *Disconnected Cluster Condition*:
      If $\tilde{w}_{\mathcal{A}} = 2$, then adding this link would result in a disconnected graph. Remove the link (i, j) from $\mathcal{O}$ *without adding it to the partial graph*.
    * Else, add the link (i, j) to the partial graph: decrement $\tilde{w}_i$ and $\tilde{w}_j$, and update $\tilde{w}_{\mathcal{A}}$ accordingly. Remove (i, j) from the ordered list $\mathcal{O}$.

Note: There is potentially a third case in which $i \in \mathcal{B}$, $j \in \mathcal{B}$, $i \neq j$; however this can only occur if there are no remaining stubs in the partial graph $\tilde{g}_{\mathcal{A}}$. This is precluded by the test for the Disconnection Condition among Type II link additions; however if the algorithm were modified to allow this, then this third case would represent the situation where graph construction continues with a new (disconnected) cluster.

Adding link (i, j) to the graph would require moving both nodes i and j from $\mathcal{B}$ to $\mathcal{A}$, decrementing $\tilde{w}_i$ and $\tilde{w}_j$, updating both $\tilde{w}_{\mathcal{A}}$ and $d_{\mathcal{B}}$ accordingly, and removing (i, j) from the ordered list O.

- Step 3 (Repeat): Return to Step 1.

Each iteration of the algorithm either adds a link from the list in O or removes it from consideration. Since there are a finite number of elements in O, the algorithm is guaranteed to terminate in a finite number of steps. Furthermore, the ordered nature of O ensures the following property.

Proposition A.3: At each point during the above construction, for any nodes $i \in \mathcal{A}$ and $j \in \mathcal{B}$, $d_i \geq d_j$.

Proof. By construction, if $i \in \mathcal{A}$ and $j \in \mathcal{B}$, then for some previously added node $k \in \mathcal{A}$, it must have been the case that $d_k d_i \geq d_k d_j$. Since $d_k > 0$, it follows that $d_i \geq d_j$. □

A less obvious feature of this construction is whether or not the algorithm returns a simple connected graph satisfying degree sequence D (if one exists). While this remains an open question, we show that if the Tree Condition is ever reached, then the algorithm is guaranteed to return a graph satisfying the intended degree sequence.

Proposition A.4: (Tree Construction). Given a graphic sequence D, if at *any* point during the above algorithm the Tree Condition is satisfied, then

(a) the Tree Condition will remain satisfied through all intermediate construction, and

(b) the final graph will exactly satisfy the intended degree sequence.

Proof. To show part (a), assume that $d_{\mathcal{B}} = 2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}$ and observe that as a result only a link satisfying Type I can be added next by our algorithm. Thus, the next link (i, j) to be added will have $i \in \mathcal{A}$ and $j \in \mathcal{B}$, and in doing so we will move node j from the working set $\mathcal{B}$ to $\mathcal{A}$. As a result of this update, we will have $\Delta d_{\mathcal{B}} = -d_j$, $\Delta|\mathcal{B}| = -1$, and $\Delta\tilde{w}_{\mathcal{A}} = d_j - 2$. Thus, we have updated the following values

$$\begin{aligned} d'_{\mathcal{B}} &\equiv d_{\mathcal{B}} + \Delta d_{\mathcal{B}} \\ &= d_{\mathcal{B}} - d_j, \end{aligned}$$

and

$$\begin{aligned} 2|\mathcal{B}'| - \tilde{w}'_{\mathcal{A}} &\equiv 2(|\mathcal{B}| + \Delta|\mathcal{B}|) - (\tilde{w}_{\mathcal{A}} + \Delta\tilde{w}_{\mathcal{A}}) \\ &= 2(|\mathcal{B}| - 1) - (\tilde{w}_{\mathcal{A}} + d_j - 2) \\ &= 2|\mathcal{B}| - \tilde{w}_{\mathcal{A}} - d_j \\ &= d_{\mathcal{B}} - d_j. \end{aligned}$$

Thus, $d'_{\mathcal{B}} = 2|\mathcal{B}'| - \tilde{w}'_{\mathcal{A}}$, and the Tree Condition will continue to hold after the addition of each subsequent Type I link (i, j).

To show part (b), observe that after $|\mathcal{B}|$ Type I link additions (each of which results in $\Delta|\mathcal{B}| = -1$) the set $\mathcal{B}$ will be empty, thereby implying also that $d_{\mathcal{B}} = 0$. Since the relationship $d_{\mathcal{B}} = 2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}$ continues to hold after each Type I link addition, then it must be that $|\mathcal{B}| = 0$ and $d_{\mathcal{B}} = 0$ collectively imply $\tilde{w}_{\mathcal{A}} = 0$. Furthermore, since $\tilde{w}_{\mathcal{A}} = \sum_{i\in\mathcal{A}} \tilde{w}_i$ and $\tilde{w}_i = d_i - \tilde{d}_i \geq 0$ for all i, then $\tilde{w}_i = 0$ for all i, and the degree sequence is satisfied. □

An important question is under what conditions the Tree Condition is met during the construction process. Rewriting this condition as $d_{\mathcal{B}} - [2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}] = 0$, observe that when the algorithm is initialized in Step 0, we have $d_{\mathcal{B}} = \sum_{i=2}^{n} d_i$, $\tilde{w}_{\mathcal{A}} = d_1$ and that $|\mathcal{B}| = n - 1$. This implies that after initialization, we have

$$d_{\mathcal{B}} - [2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}] = \sum_{i=2}^{n} d_i - 2|\mathcal{B}| + d_1 = \sum_{i=1}^{n} d_i - 2(n-1).$$

Note that minimal connectivity among n nodes is achieved by a tree having total degree $\sum_{i=1}^{n} d_i = 2(n-1)$, and this corresponds to the case when the Tree Condition is met at

initialization. However, if the sequence D is graphical and the Tree Condition is not met at initialization, then $d_{\mathcal{B}} - [2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}] = 2z > 0$, where $z = (\sum_{i=1}^{n} d_i/2) - (n-1)$ is the number of "extra" links above what a tree would require. Assuming $z > 0$, consider the outcome of subsequent Link Addition operations, as defined in Step 2:

- As already noted, when a Type I connection is made (thus adding a new node j to the graph), we have $\Delta d_{\mathcal{B}} = -d_j$, $\Delta \tilde{w}_{\mathcal{A}} = d_j - 2$, and $\Delta|\mathcal{B}| = -1$, which in turn means that Type I connections result in $\Delta\left(d_{\mathcal{B}} - [2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}]\right) = 0$.
- Accordingly, when a Type II connection is made between two stubs in $\mathcal{A}$, we have $\Delta \tilde{w}_{\mathcal{A}} = -2$, and both $|\mathcal{B}|$ and $d_{\mathcal{B}}$ remain unchanged. Thus, $\Delta\left(d_{\mathcal{B}} - [2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}]\right) = -2$.

So if $d_{\mathcal{B}} - [2|\mathcal{B}| - \tilde{w}_{\mathcal{A}}] = 2z > 0$, then subsequent link additions will cause this value to either decrease by 2 or remain unchanged, or in other words, adding additional links can only bring the algorithm closer to the Tree Condition. Nonetheless, our algorithm is *not* guaranteed to reach the Tree Condition for all graphic sequences D (i.e., we have not proved this), although we have not found any counterexamples in which the algorithm fails to achieve the desired degree sequence. If that were to happen, however, the algorithm would terminate with $\tilde{w}_i > 0$ for some node $i \in \mathcal{A}$, even though $|\mathcal{B}| = 0$. Nonetheless, in the case where the graph resulting from our construction does satisfy the intended degree sequence D, we can prove that it is indeed the $s_{\max}$ graph.

Proposition A.5: (General Construction). If the graph g resulting from our algorithm is a connected, simple graph satisfying the intended degree sequence D, then this graph is the $s_{\max}$ graph of $G(D)$.

Proof. Observe that, in order to satisfy the degree sequence D, the graph g contains a total of $l = \sum_{i=1}^{n} d_i/2$ links from the ordered list O. Since elements of O are ordered by decreasing weight $d_i d_j$, it is obvious that, in the absence of constraints that require the final graph to be connected or satisfy the sequence D, a graph containing the first l elements of O will maximize $\sum_{(i,j)\in\mathcal{E}} d_i d_j$. However, in order to ensure that g is an element of the space $G(D)$, when selecting the l links it is usually necessary to "skip" some elements of O, and conditions A.1 and A.2 identify two simple situations where skipping a potential link is

required. While skipping links under other conditions may be necessary to guarantee that the resulting graph satisfies D (indeed, the current algorithm is not guaranteed to do this), our argument is that *if these are the only conditions* under which elements of $\mathcal{O}$ have been skipped during construction *and* the resulting graph does satisfy D, then the resulting graph maximizes $s(g)$.

To see this more clearly, consider a second graph $\tilde{g} \neq g$ also constructed from the ordered list $\mathcal{O}$. Let $\mathcal{E} \subset \mathcal{O}$ be the (ordered) list of links in the graph g, and let $\tilde{\mathcal{E}} \subset \mathcal{O}$ be the (ordered) list of links in the graph $\tilde{g}$. Assume that these two lists differ by only a single element, namely $e \in \mathcal{E}, e \notin \tilde{\mathcal{E}}$ and $\tilde{e} \notin \mathcal{E}, \tilde{e} \in \tilde{\mathcal{E}}$, where $\mathcal{E} \backslash e = \tilde{\mathcal{E}} \backslash \tilde{e}$. By definition, both e and $\tilde{e}$ are elements of $\mathcal{O}$, and there are two possible cases for their relative position within this ordered list (here, we use the notation "$\prec$" to mean "proceeds in order").

- If $e \prec \tilde{e}$, then $\tilde{g}$ uses in place of e a link that occurs "later" in the sequence $\mathcal{O}$. However, since $\mathcal{O}$ is ordered by weight, using $\tilde{e}$ cannot result in a higher value for $s(\tilde{g})$.

- If $\tilde{e} \prec e$, then $\tilde{g}$ uses in place of e a link that occurs "earlier" in the sequence $\mathcal{O}$—one that had been "skipped" in the construction of g. However, the "skipped" elements of $\mathcal{O}$ will correspond to instances of Conditions A.1 and A.2, and using them must necessarily result in a graph $\tilde{g} \notin G(D)$ because it is either disconnected or because its degree sequence does not satisfy D.

Thus, for any other graph $\tilde{g}$, it must be the case that either $s(\tilde{g}) \leq s(g)$ or $\tilde{g} \notin G(D)$, and therefore we have shown that g is the $s_{\max}$ graph. □

A.2.2 Among Connected, Acyclic Graphs

In the special case when $\sum_{i=1}^{n} d_i = 2(n-1)$, there exists only one type of graph structure that will connect all n nodes, namely an acyclic graph (i.e., a tree). All connected acyclic graphs are necessarily simple. Because acyclic graphs are a special case of elements in $G(D)$, generating $s_{\max}$ trees is achieved by making the appropriate Type I connections in the aforementioned algorithm. In effect, this construction is essentially a type of deterministic

preferential attachment, one in which we iterate through all nodes in the ordered list D and attach each to the highest-degree node with a remaining stub.

In the case of trees, the arguments underlying the $s_{\max}$ proof can be made more precise. Observe that the incremental construction of a tree is equivalent to choosing for each node in $\mathcal{B}$ the single node in $\mathcal{A}$ to which it becomes attached. Consider the choices available for connecting two nodes $k, m \in \mathcal{B}$ to nodes $i, j \in \mathcal{A}$ where $d_i \geq d_j$, $d_k \geq d_m$, and observe that $d_i d_k + d_i d_m \geq d_i d_k + d_j d_m \geq d_j d_k + d_i d_m \geq d_j d_k + d_j d_m$, where second inequality follows from Proposition 3 while the first and last inequalities are by assumption. There are two cases of interest. First, if $\tilde{w}_i > 1$ and $\tilde{w}_j \geq 1$, then it is clear that it is optimal to connect *both* nodes $k, m \in \mathcal{B}$ to node $i \in \mathcal{A}$. Second, if $\tilde{w}_i = 1$ and $\tilde{w}_j \geq 1$, then it is clear that it is optimal to connect $k \in \mathcal{B}$ to $i \in \mathcal{A}$ and $m \in \mathcal{B}$ to $j \in \mathcal{A}$. All other scenarios can be decomposed into these two cases, thus proving that the algorithm's incremental construction for a tree is guaranteed to result in the $s_{\max}$ graph.

There are many important properties of $s_{\max}$ trees that are discussed in [70].

A.3 When $r_{\min} = -1$

In order to see when a degree sequence D can achieve $r(g) = -1$, we introduce a simplified version of *Cauchy-Schwarz-Burnyakovskii inequality*, which states that for any vector $\{b_1, b_2, \ldots, b_n\}$, it must be that

$$\sum_{i=1}^{n} b_i^2 \geq \frac{1}{n}\left(\sum_{i=1}^{n} b_i\right)^2,$$

with the equality holding if and only if $b_1 = b_2 = \cdots = b_n$.

Applying this inequality to a graph with l links, it follows that

$$\sum_{(i,j)\in\mathcal{E}} (d_i + d_j)^2 \quad \geq \quad \frac{1}{l}\left(\sum_{(i,j)\in\mathcal{E}} (d_i + d_j)\right)^2.$$

Expanding the squared term on the LHS and dividing both sides by 2, we have from rela-

tions (4.14-4.15) that

$$\begin{aligned} \sum_{(i,j)\in\mathcal{E}} 2d_i d_j/2 + \sum_{(i,j)\in\mathcal{E}} (d_i^2 + d_j^2)/2 &\geq \frac{1}{2l}\left(\sum_{(i,j)\in\mathcal{E}} (d_i + d_j)\right)^2 \\ s(g) + s_{\max}^{\mathcal{G}(D)} &\geq 2\, s(g_c) \\ \frac{s(g) - s(g_c)}{s_{\max}^{\mathcal{G}(D)} - s(g_c)} &\geq -1, \end{aligned}$$

which is simply another way of showing that $r(g) \geq -1$, but it proves that $r(g) = -1$ if and only if $d_i + d_j = d$ (a constant) for all $(i, j) \in \mathcal{E}$.

Recall that within $\mathcal{G}(D)$ one has $s_{\min} = Z + \hat{Z}$ as defined by (4.9), and thus this $s_{\min}$ graph corresponds to $r = -1$ if and only if for each element k one has $z_k + \hat{z}_k = z$ (a constant).

Bibliography

[1] Abilene Network. Detailed information about the objectives, organization and development of the Abilene network are available from .

[2] CENIC: Corporation for Education Network Intitiatives in California. http://www.cenic.org.

[3] Internet2 netflow: Weekly reports. Internet2 Consortium, http://netflow.internet2.edu/weekly/.

[4] Skitter, Cooperative Association for Internet Data Analysis (CAIDA). http://www.caida.org/tools/measurement/skitter/.

[5] The Rearrangement Inequality. http://en.wikipedia.org/wiki/Rearrangement_inequality.

[6] State of Washington Master Contract for Cisco Products. http://techmall.dis.wa.gov/master_contracts/intranet/routers_switches.asp, 2002.

[7] J. Abello, A. Buchsbaum, and J. Westbrook. A functional approach to external graph algorithms. *Lect. Notes Comute. Sci.*, 1461:332–343, 1998.

[8] W. Aiello, F. Chung, and L. Lu. *Random evolution in massive graphs*. Handbook on Massive Data Sets, 2001.

[9] W. Aiello, F. Chung, and L. Lu. A random graph model for massive graphs. In *Proceeding of STOC*, 2003.

[10] R. Albert and A-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74, January 2002.

[11] R. Albert, H. Jeong, and A-L. Barabási. Error and attack tolerance in complex networks. *Nature*, 406:378, 2000.

[12] D. Alderson. Technological and economic drivers and constraints in the internet's "last mile". Tech Report CIT-CDS-04-004 California Institute of Technology, 2004.

[13] D. Alderson, J. C. Doyle, R. Govindan, and W. Willinger. Toward an optimization-driven framework for designing and generating realistic internet topologies. *ACM SIGCOMM Computer Communications Review*, 2003.

[14] D. Alderson and W. Willinger. A contrasting look at self-organization in the internet and next-generation communication networks. *IEEE Communications Magazine*, July 2005.

[15] L. A. N. Amaral, A. Scala, M. Barthelemy, and H. E. Stanley. Classes of small-world networks. *Proceedings of the National Academy of Sciences*, 97:11149–11152, 2000.

[16] P. Bak. *How Nature Works: The Science of Self-organized Criticality*. Copernicus, 1996.

[17] P. Ball. Networking can be hazardous. *Nature Science Update*, March 2001.

[18] P. Ball. *Critical Mass: How One Thing Leads to Another*. Farrar, Straus and Giroux, 2004.

[19] A.-L. Barabási. *Linked: The New Science of Networks*. Perseus Publishing, 2002.

[20] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[21] A.-L. Barabási, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks. *Physica A*, 272:173–187, 1999.

[22] A.-L. Barabási, Z. Dezso, E. Ravasz, S.H.Yook, and Z. Oltvai. Scale-free and hierarchical structures in complex networks. *AIP Conference Proceedings*, 2003.

[23] A.-L. Barabási, E. Ravasz, and T. Vicsek. Deterministic scale-free networks. *Physica A*.

[24] E. A. Bender and E. R. Canfield. The asymptotic number of labeled graphs with given degree sequences. *Journal of Combinatorial Theory A*, 24:296–307, 1978.

[25] G. Bianconi and A.-L. Barabási. Competition and multiscaling in evolving networks. *Europhys. Lett.*, 54:436–442, 2001.

[26] B. Bollobás and O. Riordan. Robustness and vulnerability of scale-free random graphs. *Internet Math*, 1:1–35, 2003.

[27] S. Bornholdt and H.G. Schuster. *Handbook of Graphs and Networks: From the Genome to the Internet*. Wiley-VCH, 2003.

[28] L. Briesemeister, P. Lincoln, and P. Porras. Epidemic profiles and defense of scale-free networks. In *ACM Workshop on Rapid Malcode (WORM)*, 2003.

[29] A. Broido and k. Claffy. Internet topology: Connectivity of ip graphs. In *Proceeding of SPIE ITCom WWW Conference*, 2001.

[30] T. Bu and D. Towsley. On distinguishing between internet power-law topology generators. In *IEEE INFOCOM*, 2002.

[31] M. Buchanan. *Ubiquity: The Science of History... or Why the World Is Simpler Than We Think*. Crown, 2001.

[32] M. Catanzaro, M. Boguná, and R. Pastor-Satorras. Generation of uncorrelated random scale-free networks. *Physical Review E*, 71(027103), 2005.

[33] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Towards capturing representative AS-Level internet topologies. In *ACM SIGMETRICS*, 2002.

[34] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power-laws in internet topologies revisited. In *Procedings of IEEE INFOCOM*, 2002.

[35] F. Chung and L. Lu. The average distance in a random graph with given expected degrees. *Internet Mathematics*, 1:91–113, 2003.

[36] D. D. Clark. The design philosophy of the darpa internet protocols. *ACM Computer Communication Reviews*, 18:106–114, 1988.

[37] V. Colizza, A. Flammini, M. A. Serrano, and A. Vespignani. Detecting rich-club ordering in complex networks. *Nature Physics*, 2:110–115, 2006.

[38] Supratim Deb and R. Srikant. Congestion control for fair resource allocation in networks with multicast flows. *IEEE/ACM Transactions on Networking*, 12(2):274–285, 2004.

[39] Z. Dezsö and A.-L. Barabási. Halting viruses in scale-free networks. *Physical Review E*, 2002.

[40] M. B. Doar. A better model for generating test networks. In *Proceedings of GLOBECOM*, 1996.

[41] S. N. Dorogovtsev. Clustering of correlated networks. *Physical Review E*, 69(027104), 2004.

[42] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks. cond-mat/0106144, 2001.

[43] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press, Oxford, UK, 2003.

[44] J.C. Doyle, D. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. The “robust yet fragile” nature of the internet. *Proceedings of the National Academy of Sciences USA*, 102(41):14497–14502, October 2005.

[45] R. Durrett. *Essentials of Stochastic Process*. Springer, New York, 1999.

[46] P. Erdös and T. Gallai. Graphs with prescribed degrees of vertices. *Mat. Lapok (Hungarian)*, 11:264–274, 1960.

[47] P. Erdös and A. Renyi. On random graphs I. *Publications Mathematics*, 9:290–297, 1959.

[48] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM*, 1999.

[49] I. Farkas, I. Dereyi, G. Palla, and T. Vicsek. Equilibrium statistical mechanics of network structures. *Lecture Notes in Physics*, 650:163–187, 2004.

[50] T. Feder, A. Guetz, M. Mihail, and A. Saberi. A local switch markov chain on given degree graphs. In *FOCS*, 2006.

[51] W. Feller. *An Introduction to Probability Theory and Its Applications*. John Wiley & Sons. New York, 1971.

[52] L. Gao. On inferring autonomous system relationships in the internet. In *Proceedings of IEEE Global Internet Symposium*, 2000.

[53] Naveen Garg and Jochen Konemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, pages 300–309, 1998.

[54] C. Gkantsidis, M. Mihail, and A. Saberi. Conductance and congestion in power-law graphs. In *ACM Sigmetrics*, 2003.

[55] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *IEEE INFOCOM*, 2000.

[56] S. Jamin S. Shenker H. Tangmunarunkit, R. Govindan and W. Willinger. Network topology generators: Degree-based vs. structural. In *Proceedings of ACM SIGCOMM*, 2002.

[57] M. Hall. *Combinatorial Theory*. Blaisdell Publishing Company.

[58] C. Hedrick. Routing information protocol. RFC 1058, Internet Engineering Task Force, 1988.

[59] H.Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A-L Barabasi. The large-scale organization of metabolic networks. *Nature*, 407:651, 2000.

[60] S. Itzkovitz, R. Levit, N. Kashtan, R. Milo, M. Itzkovitz, and U. Alon. Multi-level coarse-graining of complex networks. Preprint, 2004.

[61] S. Itzkovitz, R. Levitt, N. Kashtan, R. Milo, M. Itzkovitz, and U. Alon. Coarse-graining and self-dissimilarity of complex networks. arXiv:q-bio.MN/0405011 v1, 2004.

[62] H. Jeong, Z. Neda, , and A-L. Barabási. Measuring preferential attachment in evolving networks. *Europhys. Letter*, 61:567–572, 2003.

[63] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR443-00, Department of EECS, University of Michigan, 2000.

[64] Koushik Kar, Saswati Sarkar, and Leandros Tassiulas. Optimization based rate control for multirate multicast sessions. In *Proceedings of IEEE Infocom*, pages 123–132, April 2001.

[65] Frank P. Kelly. Fairness and stability of end-to-end congestion control. *European Journal of Control*, 9:159–176, 2003.

[66] Frank P. Kelly, Aman Maulloo, and David Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of Operations Research Society*, 49(3):237–252, March 1998.

[67] P. L. Krapivsky and S. Redner. Organization of growing random networks. *Physical Review E*, 63(066123), 2001.

[68] S. Kunniyur and R. Srikant. End-to-end congestion control: utility functions, random losses and ECN marks. *IEEE/ACM Transactions on Networking*, 11(5):689 – 702, October 2003.

[69] J. Laprie. Dependable computing and fault tolerance: Concepts and terminology. In *in 12th International Conference on Fault Tolerant Computing Systems*, 1985.

[70] L. Li, D. Alderson, J. C. Doyle, and W. Willinger. Towards a theory of scale-free graphs: Definition, properties, and implications. *Internet Mathematics*, 2(4):431–523, 2006.

[71] L. Li, D. Alderson, W. Willinger, and J. C. Doyle. A first-principles approach to understanding the internet's router-level topology. In *Proceedings of ACM SIGCOMM*, Portland, OR, 2004.

[72] L. Li, V. Bahl, R. Chandra, S. Kandula, D. Maltz, and M. Zhang. Detecting user-perceived faults from packet traces in enterprise networks. Microsoft Technical Report, 2006.

[73] Steven H. Low. A duality model of TCP and queue management algorithms. *IEEE/ACM Transactions on Networking*, 11(4):525–536, August 2003.

[74] Steven H. Low and David E. Lapsley. Optimization flow control I: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, 1999.

[75] Steven H. Low, Fernando Paganini, and John C. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 22(1):28–43, Feb. 2002.

[76] Steven H. Low, Larry Peterson, and Limin Wang. Understanding Vegas: a duality model. *Journal of ACM*, 49(2):207–235, March 2002.

[77] S. E. Luria and M. Delbrück. Mutations of bacteria from virus sensitivity to virus resistance. *Genetics*, 28:491–511, 1943.

[78] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. In *Proceedings of ACM SIGCOMM*, Pisa, Italy, Sept 2006.

[79] P. Mahlmann and C. Schindelhauer. Peer-to-peer networks based on random transformations of connected regular undirected graphs. In *SPAA*, 2005.

[80] B. B. Mandelbrot. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. Springer-Verlag, New York, 1997.

[81] S. Maslov and K. Sneppen. Specificity and stability in topology of protein networks. *Science*, 296:910–913, 2002.

[82] S. Maslov, K. Sneppen, and A. Zalianyzk. Detection of topological patterns in complex networks: correlation profile of the internet. *Physica A*, 333:529–540, 2004.

[83] L. Massoulie and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Transactions on Networking*, 10(3):320–328, June 2002.

[84] A. Medina, A. Lakhina, I. Matta, and J. Byers. Brite: An approach to universal topology generation. In *Proceedings of MASCOTS*, 2001.

[85] A. Medina, I. Matta, and J. Byers. On the origin of power-laws in internet topologies. In *ACM SIGCOMM Computer Communications Review*, 2000.

[86] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: Simple building blocks of complex networks. *Science*, page 298, 2002.

[87] M. Mitzenmacher. A brief history of generative models for power-law and lognormal distributions. *Internet Mathematics*, 1:226–251, 2004.

[88] Jeonghoon Mo and Jean Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, October 2000.

[89] M. Molloy and B. Reed. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms*, 6:161–179, 1995.

[90] J. M. Montoya and R. V. Sole. Small world patterns in food webs. *Journal of Theoretical Biology*, 214(3):405–412, 2002.

[91] J. Moy. OSPF version 2. RFC 2328, Internet Engineering Task Force, 1998.

[92] M. Bogu ná, R. Pastor-Satorras, and A. Vespignani. Cut-offs and finite size effects in scale-free networks. *Eureup Physics Journal B*, 38:204–209, 2004.

[93] M. E. J. Newman. Assortative mixing in networks. *Physical Review E*, 89(208701), 2002.

[94] M. E. J. Newman. Mixing patterns in networks. *Physical Review E*, 67(026126), 2003.

[95] M. E. J. Newman. Random graphs as models of networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook of Graphs and Networks*. Wiley-VCH, Berlin, 2003.

[96] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

[97] M. E. J. Newman. Power-laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 2005.

[98] M. E. J. Newman and J. Park. Why social networks are different from other types of networks. *Physical Review E*, 68(036122), 2003.

[99] A. M. Odlyzko. Internet traffic growth: Sources and implications. In *Optical Transmission Systems and Equipment for WDM Networking II*, 2003.

[100] C. R. Palmer and J. G. Steffan. Generating network topologies that obey power-laws. In *Proceedings of GLOBECOM*, 2000.

[101] A. Papachristodoulou, L. Li, and J. C. Doyle. Methodological frameworks for large-scale network analysis and design. *Computer Communications Review*, 2004.

[102] J. Park and M. E. J. Newman. The origin of degree correlations in the internet and other networks. *Physical Review E*, 68(026112), 2003.

[103] J. Park and M. E. J. Newman. The statistical mechanics of networks. *Physical Review E*, 70(066117), 2004.

[104] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical Review Letters*, 86(14):3200–3203, 2001.

[105] R. Pastor-Satorras and A. Vespignani. *Evolution and Structure of the Internet : A Statistical Physics Approach*. Cambridge University Press, Cambridge, UK, 2004.

[106] K. Patch. Net inherently virus prone. *Technology Research News*, 2001.

[107] S. L. Pimm, J. H. Lawton, and J. E. Cohen. Food web patterns and their consequences. *Nature*, 350:669–674, 1991.

[108] D. J. Price. Networks of science papers. *Science*, 149:510–515, 1965.

[109] G. Samorodnitsky and M.S. Taqqu. *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*. Chapman and Hall, New York - London, 1994.

[110] M. A. Serrano and M. Boguná. Tuning clustering in random networks with arbitrary degree distributions. *Physical Review E*, 72(036133), 2005.

[111] M. A. Serrano and M. Boguná. Percolation and epidemic thresholds in clustered networks. cond-mat/06033353, 2006.

[112] H. A. Simon. On a class of skew distribution functions. *Biometrika*, 42:425–440, 1955.

[113] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP Topologies with Rocketfuel. *IEEE/ACM Transactions on Networking*, 12:2–16, 2004.

[114] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.

[115] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the internet hierarchy from multiple vantage points. In *Proceedings of IEEE INFOCOM*, 2002.

[116] R. Tanaka. Scale-rich metabolic networks. *Physical Review Letters*, 94(168101), 2005.

[117] J. H. van Lint and R. M. Wilson. *A Course in Combinatorics*. Cambridge University Press, 1992.

[118] F. Viger and M. Latapy. Efficient and simple generation of random simple connected graphs with prescribed degree sequence. In *Proceedings of the 11th Conference of Computing & Combinatoric*, 2005.

[119] J. Wang, L. Li, S. Low, and J. C. Doyle. Cross-layer optimization in TCP/IP networks. *IEEE/ACM Transactions on Networking*, 13, 2005.

[120] D. J. Watts and S. H. Strogatz. Collective dynamics of "small world" networks. *Nature*, 393, 1998.

[121] B. M. Waxman. Routing of multipoint connections. *IEEE Journal of Selected Areas in Communications*, 6, 1988.

[122] D. E. Whitney and D. Alderson. Are technological and social networks really different? In *Proceedings of ICCS*, 2006.

[123] W. Willinger, D. Alderson, J. C. Doyle, and L. Li. More "normal" than normal: Scaling distributions and complex systems. In *Proceedings of the 2004 Winter Simulation Conference*, 2004.

[124] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker. Scaling phenomena in the internet: Critically examining criticality. *Proceedings of the National Academy of Sciences*, 2002:2573–2580, 99.

[125] S. H. Yook, H. Jeong, and A.-L. Barabási. Modeling the internet's large-scale topology. *Proceedings of the National Academy of Sciences*, 99:13382–13386, 2002.

[126] G. Yule. A mathematical theory of evolution based on the conclusions of Dr. J. C. Willis. *F.R.S. Philosophical Transactions of the Royal Society of London (Series B)*, 213:21–87, 1925.

[127] E. Zegura, K.L. Calvert, and M.J. Donahoo. A quantitative comparison of graph-based models for internet topology. *IEEE/ACM Transactions on Networking*, 5(6), 1997.

[128] Y. Zhang, M. Roughan, C. Lund, and D. Donoh. An information-theoretic approach to traffic matrix estimation. In *Proceedings of ACM SIGCOMM*, 2003.