Topological Characteristics of IP Networks

Hamed Haddadi

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of the
University College London.

Department of Electronic & Electrical Engineering
University College London

November 11, 2008

Examination Committee:

Dr Christos Gkantsidis, Microsoft Research Cambridge

Professor George Pavlou, University College London

Dr Miguel Rio, University College London

I, Hamed Haddadi, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Department of Electronic & Electrical Engineering
University College London

Abstract

Topological analysis of the Internet is needed for developments on network planning, optimal routing algorithms, failure detection measures, and understanding business models. Accurate measurement, inference and modelling techniques are fundamental to Internet topology research. A requirement towards achieving such goals is the measurements of network topologies at different levels of granularity. In this work, I start by studying techniques for inferring, modelling, and generating Internet topologies at both the router and administrative levels. I also compare the mathematical models that are used to characterise various topologies and the generation tools based on them.

Many topological models have been proposed to generate Internet Autonomous System(AS) topologies. I use an extensive set of measures and innovative methodologies to compare AS topology generation models with several observed AS topologies. This analysis shows that the existing AS topology generation models fail to capture important characteristics, such as the complexity of the local interconnection structure between ASes. Furthermore, I use routing data from multiple vantage points to show that using additional measurement points significantly affect our observations about local structural properties, such as clustering and node centrality. Degree-based properties, however, are not notably affected by additional measurements locations. The shortcomings of AS topology generation models stems from an underestimation of the complexity of the connectivity in the Internet and biases of measurement techniques.

An increasing number of synthetic topology generators are available, each claiming to produce representative Internet topologies. Every generator has its own parameters, allowing the user to generate topologies with different characteristics. However, there exist no clear guidelines on tuning the value of these parameters in order to obtain a topology with specific characteristics. I propose a method which allows optimal parameters of a model to be estimated for a given target topology. The optimisation is performed using the weighted spectral distribution metric, which simultaneously takes into account many the properties of a graph.

In order to understand the dynamics of the Internet, I study the evolution of the AS topology over a period of seven years. To understand the structural changes in the topology, I use the weighted spectral distribution as this metric reveals differences in the hierarchical structure of two graphs. The results indicate that the Internet is changing from a strongly customer-provider oriented, disassortative network, to a soft-hierarchical, peering-oriented, assortative network. This change is indicative of evolving business relationships amongst organisations.

Publications

Journals

- R. G. Clegg, M. S. Withall, A. W. Moore, I. W. Phillips, D. J. Parish, M. Rio, R. Landa, H. Haddadi, K. Kyriakopoulos, J. Auge and D. Salmon, “Challenges in the capture and dissemination of measurements from high-speed networks”, IET Communications, Special Issue on Simulation, Analysis and Measurement of Broadband Network Traffic, 2009.

- Hamed Haddadi, Gianluca Iannaccone, Andrew Moore, Richard Mortier, Miguel Rio, “Network Topologies: Inference, Modelling and Generation”, IEEE Communications Surveys and Tutorials, Volume 10, Number 2, 2008.

- Hamed Haddadi, Steve Uhlig, Andrew Moore, Richard Mortier, and Miguel Rio, “Modeling Internet Topology Dynamics”, ACM SIGCOMM Computer Communication Review, Volume 38, Number 2, April 2008.

Conferences

- Hamed Haddadi, Raul Landa, Andrew Moore, Saleem Bhatti, Miguel Rio, Xianhui Che, “Revisiting the Issues On Netflow Sample and Export Performance”, CHINACOM 2008, Third International Conference on Communications and Networking in China, August 2008, Hangzhou, China.

- Richard G. Clegg, Raul Landa, Hamed Haddadi, Andrew Moore, Miguel Rio, “Techniques for flow inversion on sampled data”, 11th IEEE INFOCOM Global Internet Symposium, April 2008, Phoenix AZ, USA.

Workshops

- Hamed Haddadi, Damien Fay, Steve Uhlig, Andrew Moore, Richard Mortier, Almerima Jamakovic, Miguel Rio, “Tuning Topology Generators Using Spectral Distributions”, SPEC International Performance Evaluation Workshop 2008, Springer LNCS, Volume 5119, June 2008, Darmstadt, Germany.

- Hamed Haddadi, Andrew Moore, Richard Mortier, Miguel Rio, Gianluca Iannaccone, “End-to-End Network Topology Generation”, Extended abstract, ACM SIGCOMM, August 2007, Kyoto, Japan.

- Hamed Haddadi and Lionel Sacks, “Networks Modelling Using Sampled Statistics”, Proceedings of LCS2006, London Communications Symposium: The Annual London Conference on Communication, University College London , September 2006.

- Hamed Haddadi, Lionel Sacks, “Passive Monitoring Challenges on High-Speed Switched Networks”, Proceedings of LCS2005, London Communications Symposium: The Annual London Conference on Communication, University College London , September 2005.

- Hamed Haddadi, Eduarda Mendes Rodrigues, Lionel E Sacks, “Development of a Monitoring and Measurement Platform for UKLight High-Capacity Network”, Proceedings of PREP2005, University of Lancaster, UK, April 2005.

- Hamed Haddadi, Eduarda Mendes Rodrigues, Lionel E Sacks, “Applications of GridProbe Technology for Traffic Monitoring on High-Capacity Backbone Networks, Data-Link Layer Simulation Approach”, Extended abstract, Proceedings of IEEE INFOCOM 2005, Miami, Florida, USA, March 2005.

Technical Reports

- Hamed Haddadi, Damien Fay, Almerima Jamakovic, Olaf Maennel, Andrew W. Moore, Richard Mortier, Miguel Rio, Steve Uhlig, “Beyond Node Degree: Evaluating AS Topology Models”, Technical Report UCAM-CL-TR-725, University of Cambridge, Computer Laboratory, July 2008.

- Richard G. Clegg, Hamed Haddadi, Raul Landa, Miguel Rio, “Towards Informative Statistical Flow Inversion”, May 2007, `http://arxiv.org/abs/0705.1939`.

Acknowledgements

Completing a PhD is a really exciting experience: a constant battle with interesting, and sometimes unpleasant, problems, situation changes, fight for funding, etc. It is mainly the people that you meet along the way that make the experience extremely rewarding. My PhD path had a lot of obstacles, changes of directions, supervisors, institutions and jobs. However a large number of people helped me on this path and I wish to mention just a few of these amazing people.

I wish to acknowledge my mentors, advisors, and role models, Dr Miguel Rio, Dr Andrew Moore and Dr Richard Mortier. They suffered the most from my disasters! Every time I needed help, they sincerely went out of their way to help me, especially when I was directionless and without any supervision. Words can not say how greatly I appreciate your help.

Most of this work was done through collaboration with Dr Damien Fay and Dr Steve Uhlig. They helped me in every stage of preparing this thesis, from detailed experiments to analysis, from writing code to teaching me mathematics. I enjoyed every minute of working with you, not to mention continuous hours on the table football with Damien.

I also appreciate all the help of Professor Saleem Bhatti and Dr Gianluca Iannaccone. They helped me realise what research means and opened my eyes to scientific thinking. My time with Gianluca in Intel was certainly what gave direction to my otherwise confused research career. Saleem was also there for me at the hardest times of PhD when I had to transition between supervisors and he helped me stay in control of situation.

The work in this thesis was done while I was visiting the Network and Operating Systems group at the University of Cambridge Computer Laboratory and Churchill College. I wish to acknowledge all the staff and students there. I had a great time and want to specifically thank Professor Jon Crowcroft, Dr Pan Hui, Dr Tim Griffin and Dr Eiko Yoneki for all their help and encouragement. I wish to also acknowledge my friends in the Cambridge Systems and Networking group at Microsoft Research Cambridge, specifically my examiner Dr Christos Gkantsidis; without his insight and great knowledge of networking this thesis would not have been completed. Any remaining errors are purely mine.

I am deeply grateful to all of my colleagues at my home group at UCL, the Networks and Services Research Laboratory, especially Professor George Pavlou, Dr Richard Clegg, Mr Raul Landa and Dr Eduarda Mendes Rodrigues. I have had a great time at UCL and it is mainly because of friends like you.

I wish to also acknowledge all the staff and students who I met while at Intel Research Cambridge, especially my dear friends Dr Vivek Mhatre, Mr Christian Schwarzer and Dr Valeria Baiamonte. I also am thankful to staff and students at the University of Adelaide in the School of Mathematical Sciences. I had a great experience there and learnt a lot about the wonders of BGP. Thanks to Dr Matthew Roughan for arranging my visit, and many thanks to my great friends, Dr Olaf Maennel, Dr Belinda Chiera and Mr Ashley Flavel.

I wish to also thank all the partners in the UKLight MASTS project and my funding body, the EPSRC, for enabling me pursue this career path.

Last, but most definitely not least, I wish to acknowledge all the help from my friends and family. My parents who devoted their lives and their comfort for my success; my family, especially Mr Reza Hashemy, were always there for me; and my friends such as Mr Hassan Moussavi, Mr Kamran Vejdani, Dr Hassan Hashimi, Dr Jonas Griem, Dr Jason Spencer, Dr Aleksandar Lazarevic, Mr Kamal Khaveripour, Mr Farhad Heravi, Mr Shu Yan Chan, and all my dear friends who are too many to list, and it would take many pages of acknowledgements! I hope I will deserve your valuable friendship all my life.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The Internet is vital to the operation of our modern society. As such, it is essential for operators to have a thorough understanding of their networks and ultimately the Internet operation, in order to meet the demands of the customers. Dynamic growth and changes in the Internet make it difficult to analyse the performance of different applications and protocols. Research in the structural properties of Internet topology is essential for studies such as failure location and fault-finding, virus propagation models, improving routing algorithms, and analysis of network growth and capacity planning strategies.

In this thesis I provide insight into Internet topology research, focusing on the organisational level topology (Chapter 2). I highlight a variety of shortcomings with current topology generators and datasets (Chapter 3), and present appropriate ways to compare datasets and design generators (Chapter 4). I demonstrate how additional network measurements can enhance the current view of the Internet topology. I also compare the properties of AS topologies relying on different sets of observations, providing insight into different aspects of the Internet topology and its evolution (Chapter 5).

1.1 Structure of the Internet

The Internet is a large, complex, decentralised and arguably self-organised network, formed of hundreds of millions of end devices such as computers, mobile phones and sensors, connected together via Internet Service Providers (ISP) and backbone connectivity providers.

From an operational point of view, the Internet is formed as a network of networks. Those constituent networks are referred to as Autonomous Systems (AS), and are often driven by self-interest economic and fiscal reasons. Although an organisation, such as a large ISP, can have multiple AS numbers in different locations, it may also be the case that a large number of organisations, such as UK higher education institutes, share an AS number. However, from the inter-AS routing perspective and traffic engineering, the AS numbers in the routing messages are an important part of the routing process. Such complications make it difficult to analyse the Internet size, its topology and the geographic location of ASes.

Today, the Internet is an essential part of international commerce, trade and culture. The Internet market is a competitive open market, with many ISPs competing for provisioning of services. Hence, the Internet is constantly undergoing changes. Edge ISPs aggressively add peering relationships with others in order to avoid paying transit charges, while larger ISPs constantly review and reconsider the peering policies with their neighbours based on the cost, utilisation and Quality of Service (QoS) agreements. Alongside business relationships, the failure of devices in networks, maintenance operations and addition of new links and routers all lead to constant changes in the network topology. Understanding these changes is important for understanding the operation of the Internet and the related applications and protocols. Hence there has been a great deal of research focused on analysing the topological characteristics of the Internet and there is need for further research. trying to characterise the dynamics of the Internet.

1.2 Motivations for Topology Research

Two decades ago one could easily obtain a complete map of the network, showing all the connections between various institutions and their respective characteristics such as bandwidth. However today the Internet is formed of about 30,000 organisations[1]. Each AS typically includes many routers and end hosts. Clearly, it is no longer possible to visualise the topology of such a large network as graphs of nodes and links, even if such a topology graph were available.

Today there are a large number of research projects focusing on Internet topology collection, such as Skitter,[2] Dimes [SS05] and RouteViews.[3] Such approaches rely on different measurement methodologies. Some rely on active probing and measurement, and some use passive collection of routing data. The availability of such rich data has increased our understanding of the Internet topology. However it has also become evident that despite all the efforts, researchers still have a limited visibility of the real Internet topology due to measurement biases. As a result, the Internet topology models which are also derived form the collected data sets tend to become biased as time goes on, as shown in Sections 3 and 5.

Performance of Internet protocols and applications is also highly related to geographic aspects and the structure of the network [GMZ03]. Treating the Internet simply as a graph of nodes and edges is not satisfactory. As I show in this thesis, our knowledge of Internet structure at core and edge is not comprehensive. In this thesis I focus on characteristics of the Internet such as dynamics of growth and connectivity amongst nodes, visibility of links and graph-related aspects.

1.3 Challenges

Lack of accurate mathematical models and topology maps of the Internet at router and AS level, despite great efforts by the research community, is due to several challenges. Internet topologies are constructed typically using passive and active measurements. Same data are used for many purposes, including construction of *realistic simulations* and analysis of business relationships between ISPs. However, the measurements, and hence the models relying on them, are subject to a large number of artifacts.

The first challenge, and arguably the most difficult one to overcome, is the inference of the actual topology. At the router level, presence of software and hardware firewalls, and traffic tunnelling services

[1]Based on assigned AS numbers; It is possible for numerous organisations to share an AS number, or an organisation having multiple AS numbers. This makes it practically impossible to obtain an accurate number of organisations forming the Internet.

[2]`http://www.caida.org/tools/measurement/skitter/`

[3]`http://www.RouteViews.org/`

have all made it difficult to obtain accurate router level topologies using active measurements. At the AS level, the routing protocols do not reveal complete information.

Lack of accurate topology data has made it difficult for scientists to accurately model the Internet topologies. Initial discoveries of the Internet topologies led to researchers modelling the Internet as scale free networks [BA99]. It was soon discovered that the inaccuracies in measurements may bias the derived models. Many links such as redundant links between routers, back up links and peering relationships between ASes where not observed by the passive and active measurement methods. More importantly, the evolution of the Internet has not been studied intensively and most previous analysis have focused on addition of nodes and links over time, rather than paying attention to the architectural dynamics. These challenges have started a whole new breadth of research in the topology modelling area in order to improve on our current understanding of the Internet.

Generation of the Internet topology calls for a model that achieves a good balance between keeping global (structural) characteristics and more local properties like node degrees and local interconnection structure. In the topology generation literature, current research focuses on distribution-driven methods, which capture some global characteristics of the topology. They rule out the randomly generated graphs and aim at attaching meta-data information (metrics) to the links and routers generated in a graph. Information about node (e.g., customer and provider) relationships, the delay and bandwidth, would also be of significant value to researchers using those graphs. Such information is currently not available and it is difficult to infer using passive and active measurement techniques.

In addition to above challenges, validation of the models is also difficult due to the constantly evolving nature of the Internet. Researchers have recently been paying attention to the evolving nature of the networks and its effects on network planning and provisioning. Another important area of research is understanding the dynamics and business incentives for addition of nodes (routers or organisations) and links between them. Gaining insight into the nature of measurement processes and their biases on the analysis of Internet topology research is another important objective of researchers. The challenges overviewed in this section are extensively studied in Section 2.2.

1.4 Contributions

The main contributions of this thesis are the following:

- Extensive analysis of currently available Internet topology generators and compare them to a wide range of observed Internet AS level topologies;

- Demonstrating the improvements in the accuracy of the structural properties of inferred topologies by additional measurements;

- Proposing a new cost function for analysis of Internet topologies: the *weighted spectral distribution*, constructed from the eigenvalues of the normalised Laplacian matrix, or graph spectrum;

- Using the proposed metric to tune parameters for a set of Internet topology generators, enabling these models to effectively match a particular measured Internet topology.

- Presenting a study of two views of the evolving Internet AS topology and expose apparent inconsistencies between these two inferred AS topologies and their evolution, highlighting structural dynamics of the Internet.

I illustrate that the core of the Internet is becoming less dominant over time, and that edges at the periphery are growing instead. I demonstrate a departure from a preferential attachment, tree-like disassortative network, toward a network that is flat, highly-interconnected, and assortative. This challenges common belief about the Internet being a scale free network, dominated by preferential attachment and incremental growth of nodes and links. The change in growth trend of the Internet calls for deeper study into business relationship models of the ISPs. In each of the chapters, I provide a detailed analysis and breakdown of the above contributions. I also expand on the impact of these contributions further in 6.1.

In this thesis I have focused on Internet AS topologies. However the measures proposed, especially the Weighted Spectral Distribution, can be used to compare other topologies such as social networks, web graphs, biological networks and router level topologies. As a future research work I am working in collaboration with social scientists and computational biologists in order to extend the uses of these methods.

1.5 Thesis Outline

The rest of this thesis is organised as follows. Chapter 2 provides an overview of the latest research in the field of network topology over the past decade. I also provide insight into the challenges involved in collecting topology data and providing realistic topology models. I bring together an analysis and summary of techniques for inference, modelling and generation of the Internet topology at router and AS level.

In Chapter 3, I perform a thorough comparison of topologies generated from several different models against a set of measured AS topologies by using a large set of topological metrics in the analysis. This analysis reveals that current topology generators fail to capture the complexity of the local interconnection structure between ASes, despite matching degree-based properties of the AS topology reasonably well. Using a collection of AS topologies from many measurement locations, I demonstrate that adding more measurement locations significantly affects local structure properties such as clustering and node centrality while not significantly affecting degree-related metrics.

When using the topology generators, I realise that a large number of generators have a number of parameters, without any guidelines on how to set these parameters to get topologies of various sizes. Chapter 4 presents the results of optimisation of the parameters of these topology generators to match a given Internet topology. The optimisation is performed either with respect to the link density, or to the spectrum of the normalised Laplacian matrix. I show that using this metric the graph properties can be better represented using most topology generators.

An important requirement of future topology generators is the ability to create dynamics models of networks that take into account the growth of networks and the failures of nodes and links. Chapter 5 illustrates the evolution of the AS topology as inferred from two different datasets over a period of seven years. I use a variety of metrics to analyse the structural changes in the Internet AS topology. The results indicate that the Internet is changing from a core-centred, strongly customer-provider oriented, disassortative network, to a soft-hierarchical, peering-oriented, assortative network. The findings indicate that traceroute-based approaches may fall short in correctly sampling the periphery of the AS topology, while the inter-domain routing dataset does not perfectly sample the inner-most core of the network. Such findings call for new efforts in the research community to devise more comprehensive measurement tools.

Finally, in Chapter 6 I summarise the contributions of the thesis, explain on the limitations of the work presented, and suggest possible directions for future research.

Chapter 2

Literature Review

Accurate measurement, inference and modelling techniques are fundamental to Internet topology research. Spatial analysis of the Internet is needed to develop network planning, optimal routing algorithms and failure detection measures. A first step towards achieving such goals is the availability of network topologies at different levels of granularity, facilitating realistic simulations of new Internet systems.

The main objective of this chapter is to familiarise the reader with research on network topology over the past decade. I study techniques for inference, modelling and generation of the Internet topology at both router and AS level. I also compare the mathematical models assigned to various topologies and the generation tools based on them.

2.1 Introduction

The Internet connects millions of computers, sensors, monitoring devices and IP telephony devices together, offering many applications and services such as the World Wide Web, email, and content distribution networks. Hosts on the Internet are connected via thousands of ISPs. An ISP contains one or more ASes depending on its size. An AS is a set of routers within a single administration domain, such as a university or corporate network.

By convention, the Internet is built upon two domain categories, *transit* and *stub*. A transit AS usually carries traffic between other domains. A stub AS, such as a corporate network, is one which has connections to end hosts and relies on at least one transit AS for connectivity to the rest of the Internet. Stub ASes usually do not enable IP packets to transit their networks, if they are not sent or received by an end host within the network. Figure 2.1 displays a simplified version of this structure.

In Figure 2.1, transit domains carry traffic between customer ASes, ISPs or Stub Domains. The ISPs may have exchange (peering) relationships among themselves for resilience and cost-saving purposes. Some ASes of ISPs are attached to more than one transit AS. This is a back-up measure increasingly being taken by corporate networks and business customers in order to ensure the existence of alternative routes to the Internet, should their main provider fail. It is also a technique for traffic engineering, allowing traffic to be sent over links of different performance. This strategy is called multi-homing and is also displayed in Figure 2.1.

The growth of the Internet and the overlay networks which rely on it has led to emerging applica-

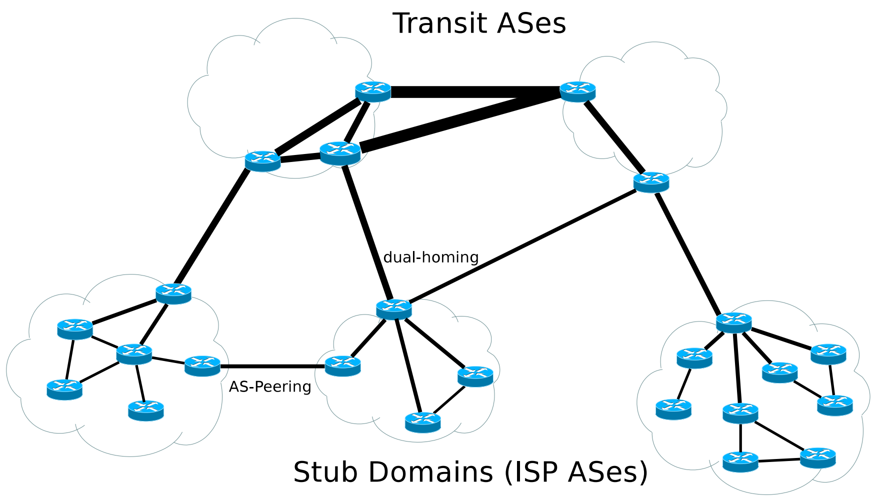

Figure 2.1: An abstract part of the Internet, link widths represent relative bandwidth.

tions and properties which have not been considered in the current topology inference and generation tools. Dynamic reconfiguration of routers and firewalls, changes in routing policies of ISPs, overlay networks, peer-to-peer networks, increasing use of Virtual Private Networks (VPNs) and protocols such as Multi-Protocol Label Switching (MPLS) and tunnelling techniques, multi-homing, on-demand circuit set-up and bandwidth allocation for home entertainment and video conferencing and the increased existence of mobile devices and laptops has caused the topology of networks to be in constant change as a result of addition, removal and reconfiguration of routers, links, devices and organisations.

In this section, I introduced the basic concepts of the Internet's operation and the need for network topology inference, modelling and generation. Section 2.2 reviews the challenges of topology inference, modelling, generation and validation. Section 2.3 describes the inference of router-level topologies of the ISPs and the AS-level topology of the Internet and the impact of geographical location of the nodes on inference techniques. Section 2.4 discusses the statistical and hierarchical models which are used to represent the topologies of the Internet at AS and router-level. In Section 2.5 I overview the tools available for topology generation. Finally in Section 2.6 I introduce possible future research directions and conclude the chapter.

2.2 Topology Research Challenges

The Internet topology is usually investigated at two levels. The Internet AS-level topology is of interest to those interested in the relationships between the networks that constitute the Internet. For example, understanding the global Internet connectivity and the business relationships between ISPs. Within an AS, the router-level topology map of ISPs is important to perform optimum network planning and to minimise the impact of router and link failures.

There are many challenges in inferring and generating realistic Internet topologies. Information on network topology, routing policies, peering relationships, resilience and capacity planning are not usually publicly available as they are considered sensitive business information by the ISPs. Instead, researchers try to infer the required data by using passive and active measurement methods to produce snapshots of the global Internet or individual ISP topologies. The fundamental problem of these techniques is the lack of ground truth of the Internet topology. Moreover, the constant evolution of the Internet leads to poor perceptions and models, as the underlying measurements are not well understood. In this section, I discuss these challenges in turn.

2.2.1 Inference of topologies

At the AS-level, it is not possible to obtain a consistent map of the actual AS-level topology of the Internet due to the constantly changing nature of the Internet topology. Operators are constantly reviewing their peering agreements. AS operators do not disclose their peering relationships and traffic exchange policies with other ASes. Connectivity between ASes is instead inferred from inter-domain routing protocols, primarily the Border Gateway Protocol (BGP) [RLH06a]. However BGP data collected from various points on the Internet is not enough to provide a user with a *complete* map of the Internet at AS-level.

Challenges also exist when trying to get the router-level topology of a single AS. The router-level topologies of ISPs are also dynamic and constantly evolving due to failures, maintenance and upgrades. Network operators are not willing to publicly release the maps of their network topology; this is sensitive information that may reveal strategic planning decisions and may also be used by attackers that may target the weak points of the network.

The most widely used tool for inference of router-level topologies is the *traceroute* tool [Mal93]. One problem with traceroute is that it is known to miss alternative links between routers. Another problem is *aliasing*. Routers have multiple interfaces with separate IP addresses. During the inference process, each of these interfaces may be reported as a different router. This problem is referred to as aliasing [SMW02]. I will discuss these issues in detail in Section 2.3.

2.2.2 Modelling the Internet

Researchers have made significant efforts to model the characteristics of the Internet. The major problem currently in this field is the absence of detailed information about inferred topologies. Many of these models are based on datasets that are known to be incomplete and prone to errors due to the nature of the collection process involved, discussed in detail in Section 2.3.

Due to above challenges, it is difficult to estimate the growth potential and characteristics of the internet. This is a vital requirement for network traffic engineering purposes. Section 2.4 describes many of the widely used models.

2.2.3 Validation of Models

Validation of generated topologies can be done by comparison to real topologies. Another common method is to compare the statistical characteristics of a generated topology with the input parameters and requirements such as certain node degree distributions or connectivity matrices. As there is no real *snapshot* of the Internet traffic or its topology, it is difficult to devise a method to benchmark the success of a topology generator or the inference of a topology, however the topologies are compared with incomplete datasets.

When inferring the router-level topology of a medium sized ISP, it may be possible to request the operator to verify the results, as done by Spring et al. [SMW02]. However, as mentioned before, operators are unlikely to reveal such information, although they may indicate the success level of an inference method as a percentage of routers or links discovered [DKF+07]. BGP and AS ownership data can also be validated by relevant Internet domain registries, although the information held by such authorities is not continuously updated and is thus often inaccurate.

2.3 Topology Inference

In this section I discuss recent efforts for inference of the AS-level topology of the Internet and router-level topology of ISPs. It is essential to note the intersection of inference with measurement. Inference-based statistics are subject to the underlying measurement process and the assumptions which have been made on the level of accuracy and details of the measurement process. Thus, inaccurate inference methods lead to unrealistic models.

Topology inference works usually fall in two categories: Router-level and AS-level. In related literature, Donnet and Friedman [DF07] also mention the IP interface and the Point-of-Presence (PoP) maps. The IP interface addresses are usually aliases for the same router and I mention the problems associated with resolving such aliases in this section. Inferring PoP level maps is a difficult task due to lack of publicly available datasets or tools. Hence they are sometimes made available by network operators, or inferred indirectly from IGP routing data.

2.3.1 ISP Router-Level Maps

In this part, I discuss the recent efforts and tools for discovering the Internet's router-level topology, also known as its *IP layer* or *layer 3* topology. These methods are usually based on the *traceroute* tool. Traceroute is the basic tool for discovering the paths that packets take in the Internet. Nearly all attempts to extract routing and topology information of the Internet at router layer use traceroute.

Traceroute works by sending multiple Internet Control Message Protocol (ICMP) [Pos81] packets with an increasing Time To Live (TTL) field in the IP header. When a packet with a TTL of one reaches a router, it discards the packet and sends an *ICMP time exceeded* packet to the sender. The traceroute tool uses the IP source address of these returning packets to produce a list of routers that the packets have traversed on their route to the destination. By incrementing the TTL value after each response, the overall path taken by the packets can be inferred.

Mercator

One of the first tools which relies on traceroute for mapping sections of an ISP is *Mercator*, introduced by Govindan et al. [GT00]. The aim of Mercator is to build a nearly complete map of the transit portion of the Internet from any location where Mercator is run, using hop-limited probing [ELR+96]. By using multiple source points, including source-route probe capable routers, it is possible to find cross links and avoid discovering only a tree-like structure. Mercator sends a UDP message to a high port number on the router and receives back an ICMP reply. If two source addresses of the reply message are the same, they are from the same router. This operation relies on the requirements for the Internet hosts as described in [Bra89]. This is a technique for resolving alias, by identifying the interfaces belonging to the same router.

The challenges faced by Mercator are due to the fact that it does not attempt to cover the whole spectrum of a network due to randomised process and the fact that many routers do not forward traceroutes for source-routing in the way that Mercator requires.

Skitter

One of the most widely used datasets is that collected by the *Skitter* project[1]. Huffaker et al. [HPM+02] state the project focus as "active measurement of the topology and round trip time (RTT) information across a wide cross-section of the Internet".

Probing uses the traceroute tool. IP addresses are then mapped into their corresponding origin AS. The disadvantage of such a tool is the large amount of data that it produces, from a number of sources currently placed in over 25 locations worldwide. This leads to the inherent problems of traceroute such as aliasing on a wider scale as multiple sources are involved. Skitter does not attempt to resolve aliases.

Rocketfuel

In an attempt similar to Mercator, Spring et al. in the Rocketfuel project [SMW02] try to infer the maps of ten ISPs, consisting of backbones, access routers and directly connected neighboring domain routers. Validation is attempted by using some of the ISP's own topology data. Direct probing techniques are used to filter the traceroutes on the ISP of interest, using BGP tables information from RouteViews. A BGP table maps destination IP address prefixes to a set of AS paths that can be used to reach that destination. Public traceroute servers are used as vantage points for the traceroutes.

Rocketfuel uses the direct probing method, as suggested by Govindan and Tangmunarunkit [GT00]. In order to ensure correct resolution of aliases, Rocketfuel also uses the `IP_ID`[2] field of the router's responses to probe packets, which is incremented by the router. The source sends two probe packets to the two interfaces that are thought to be aliases of the same router. If consecutive responses from the interfaces increment the `IP_ID` by a small value, it indicates that the same IP stack is running on the same router with multiple interfaces, hence the interfaces are believed to belong to the same router. Otherwise, the interfaces belong to two distinct routers.

Network cartographer

Another tool for inference and mapping of a network topology is the *network cartographer* (nec) mapping software introduced by Magoni and Hoerdt [MH05]. The *nec* tool is a traceroute-based mapper from multiple traceroute servers, finding routers and links and producing router-level connectivity graph. The major difference between *nec* and Rocketfuel [SMW02] is that *nec* has wider scope while Rocketfuel focuses on a single ISP. Unlike Rocketfuel, where few hosts target thousands of IP addresses, *nec* uses many traceroute webservers to a limited set of chosen IP addresses. Figure 2.2 displays the steps involved in an *nec* mapping query, sent to two traceroute servers *A* and *B*.

---

[1] `http://www.caida.org/tools/measurement/skitter/`

[2] The identification field in the IP header is used to aid in assembling the fragments of a datagram.

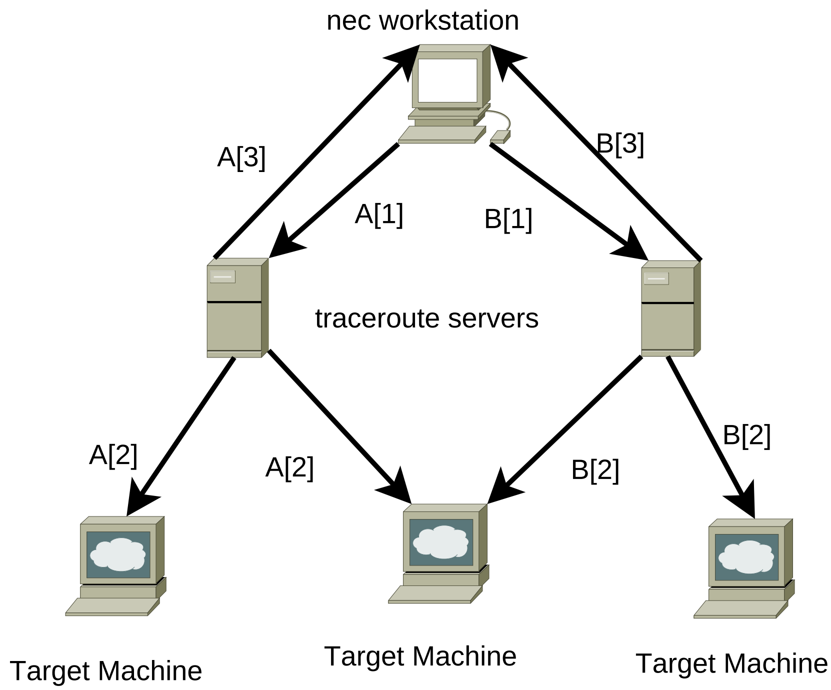

Figure 2.2: nec mapping steps, figure courtesy of [MH05]

In the first stage, the queries are sent from the workstations to the traceroute servers. In the second stage, traceroute servers query the selected IP addresses. In the final stage the results of the traceroutes are sent back to the *nec* mapping workstations.

DIMES

The DIMES project [SS05] attempts to build a router-level map of the Internet. In this project, the DIMES agent, which can be installed on any computer connected to the Internet, performs Internet measurements such as traceroute and ping at a low rate, sending the results to a central collection station at regular intervals. The advantage of the DIMES approach over previous traceroute based mapping tools is that the probing process is done across many locations in the world, giving a more complete map of the Internet router-level topology. However, due to the large number of vantage points and collection of overlapping measurements, removing the redundancies in the data is a complicated process. Moreover, DIMES also does not attempt to resolve router aliases.

2.3.2 Comparison of traceroute-based methods

In this section I have listed a number of methods for inferring router-level connectivity information. These methods have evolved over time from single source traceroute probes to universally distributed probing agents. Table 2.1 displays a summary of the characteristics of these methods.

It can be observed that the trend of inference tools has moved from single-source, static maps to those spread across many sites and constantly updating their database. It is interesting to note that there are no maintained maps with alias resolution and this may lead to incorrect assumptions about the growth of the Internet router level topology.

Table 2.1: Comparison of traceroute based methods

Tool	Released	Alias resolution	Updated	Probes
Mercator	1999	YES	NO	Single
Skitter	1999	NO	YES	Multiple
Rocketfuel	2002	YES	NO	Single
nec	2003	NO	NO	Multiple
DIMES	2004	NO	YES	Multiple

2.3.3 Accuracy of traceroute maps

Most of the work in discovering router-level topology of ISPs relies on the traceroute tool. Achlioptas et al. [ACKM05] discuss some of the problems associated with traceroute. They explore the mathematics of the sampling bias of traceroute, confirming that even when a given node degree distribution is Poisson, after traceroute sampling, the inferred node degree distribution exhibits power law properties. It is difficult to remove this bias as shown by Clauset and Moore [CM05], as the number of sources required to compensate for the bias in traceroute sampling grows linearly with the mean degree of the network.

Lakhina et al. [LBCX03] analyse the effects of such traceroute sampling techniques on random graphs and conclude that when graphs are sampled using traceroute-like methods, the resulting degree distribution can differ significantly from the the underlying graph. For example, given a sparse Erdös-Rényi random graph, the subgraph formed by a collection of shortest paths from a small set of random sources to a larger set of random destinations can exhibit a degree distribution remarkably like a power-law. The implementation of sampling in the paper is performed on the measurements from Skitter, Mercator, the dataset used by Faloutsos et al. [FFF99] and the Pansiot-Grad [PG98]. In studies of the four traces, the sampled subgraph shows differences in degree distribution and other characteristics from the original graph.

Teixeira et al. [TMSV03] look at path diversity (number of available paths) in the Sprint network[3] and ISPs explored by Rocketfuel. The Rocketfuel path diversity discovery is found to be at extreme cases, either over-estimating or finding very little diversity, again due to the use of traceroute. The differences between the Sprint data and Rocketfuel inferred maps are due to non discovery of backup links, lack of vantage points, incomplete traceroute information, path changes in a traceroute and incorrect DNS names.

Deploying a large number of monitors usually results in having to process large datasets from each monitor. Donnet et al. [DRFC06] try to find out the amount of redundancy across datasets, focusing on the CAIDA Skitter datasets. They discover that around 86% of a given monitor's probes are redundant in a sense that they visit router interfaces which have already been visited by the monitor, especially those closer to the monitoring station. It is also observed that many of the probes are redundant in a monitor's dataset as they already have been visited by the other monitors, particularly those at an intermediate distance (between 5 and 13 hops).

[3] `http://www.sprintlink.net/`

As a result of the traceroute sampling bias, there has been ongoing effort in order to modify traceroute behaviour. Augustin et al. [ACO+06] propose *Paris traceroute*, which is a modified version of traceroute with ability to discover redundant paths. One of the issues when using traceroute arises due to the Equal Cost Multi Path (ECMP) load balancing deployed by multi-homed stubs and network operators. This leads to traceroute taking different paths on each occasion as shown in Figure 2.3. Paris traceroute looks into the effects of load balancing and its frequency on traceroute anomalies. Load balancing can be done per packet, per flow or per destination IP address.

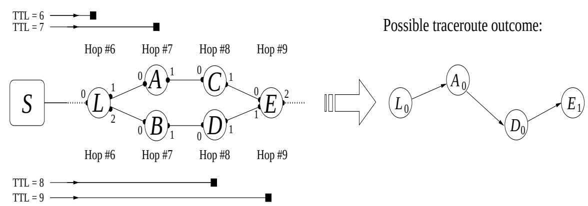

Figure 2.3: Traceroute false reporting, Figure provided by [ACO+06]

Augustin et al. show that by manipulating the ICMP sequence number and checksum in the ICMP packet header, it is possible to ensure that all the packets on traceroute take the same path. This leads to discovery of more possible routes. With this method it is also possible to report on the loops and cycles in ordinary traceroute reports. Paris traceroute is suggested as an alternative to the ordinary traceroute, rather than as a topology mapping tool, hence it does not attempt to resolve any router aliases.

Dall'Asta et al. [DAHB+06] find that the node and link detection probability depends on statistical properties of elements such as betweenness centrality. Hence the shortest path routed sampling, or sampling the network from a limited set of sources as performed by traceroute, provides a better characterisation of underlying graphs with broad distributions of connectivity, such as the Internet. The studied model analyses the efficiency of sampling in graphs with heavy-tailed connectivity distributions and looks at metrics such as the node degree distribution. The conclusion drawn is that unlike homogeneous graphs, in those with heavy-tailed degree distribution such as the Internet, major topological features are easily captured though details such as the exponent of the power laws. However this behaviour appears to suffer from biases which are result of the sampling process and affect the accuracy of results.

The studies in this section may imply that traceroute is not a suitable tool for detailed analysis of the Internet router-level maps. However it is still widely used for topology measurement and it is in reality the only scalable and available tool.

2.3.4 AS-Level Internet maps

The other important level of the Internet topology is the AS-level topology. The freedom of AS administrators to change their traffic exchange relationships with other providers has led to a constantly evolving Internet topology at router and AS level. Obtaining the AS graph can enable better design of routing algorithms and traffic engineering between various ASes.

BGP information at border routers is kept consistent by receiving BGP update messages from other ASes. BGP updates contain multiple route announcements and withdrawals. These announcements indicate that new network sections are available to the routers or a policy change is enforced to prefer an alternative path over an existing one. Withdrawals occurs when an existing route is replaced by a new route to a destination prefix by means of a withdrawal message. These messages inform the withdrawal of links and addition of new links and contain the *AS-Path* travelled by the advertisement. Each router along the path prepends its own AS number to the AS-path in the BGP message.

The AS-path is needed to avoid loops in the BGP route selection process. The AS-paths, in conjunction with the AS prefix, are also used to decide on what is the best next hop to use for sending a packet to a destination. An edge-router may not have complete view of the BGP status of the Internet and may have a default path to a tier-1 provider. Tier-1 providers have default-free BGP information so that they can forward all the packets to the correct destination. IP forwarding requires that all routers within an AS are aware of all the prefixes which are learned by the edge routers from other ASes.

Some attempts on AS-level topology discovery were based on using traceroute data. Inference of AS-level maps from traceroute data includes problems not immediately noticed. Mapping of an IP address to the correct AS number incorporates challenges which are discussed by Mao et al. [MRWK03]. They propose techniques for improving mapping of IP addresses to the corresponding ASes. These techniques rely on a measurement methodology for collecting both BGP and traceroute paths at multiple vantage points and using an initial IP-to-AS mapping derived from a large collection of BGP routing tables.

The difficulties arise due to the fact that the BGP table data and the actual path taken by packets can be inconsistent due to new route aggregation/filtering and routing anomalies [GW02]. The WHOIS data is also not always up to date due to company mergers, break ups and IP address re-allocations. An improvement can be made by collecting a large amount of information from BGP routing tables, BGP update messages and reverse DNS lookups in order to help traceroute build a more accurate AS-level map of the Internet.

Gao's seminal paper [Gao01] is one of the first attempts to present an AS graph inferred from the Oregon RouteViews BGP data. The provision of such a map has enabled classification of AS relationships into customer-provider, peering and sibling relationships. Figure 2.4 displays examples of the types of relationship between different ISPs.

A customer pays its provider for Internet connectivity and does not transit any traffic between its providers. A pair of peers agree to exchange traffic between their customers by sharing the cost of the peering links and eliminating traffic charges between each other. A pair of small ISPs may provide additional connectivity or backup connectivity to the Internet to each other in form of a sibling relationship.

Despite the presence of such contractual agreements, there is little publicly available information about inter-AS relationships. The Routing Policy Specification Language [AVG+99] can be used to register information about peering relationships but this information is not always accurately published due to its sensitive business nature. However it is possible to infer such information from the BGP routing tables. Gao proposed heuristic algorithms for such discovery, and then validated some of the results by using a Tier1 ISP's internal information. The discovery of the relationships is based on the

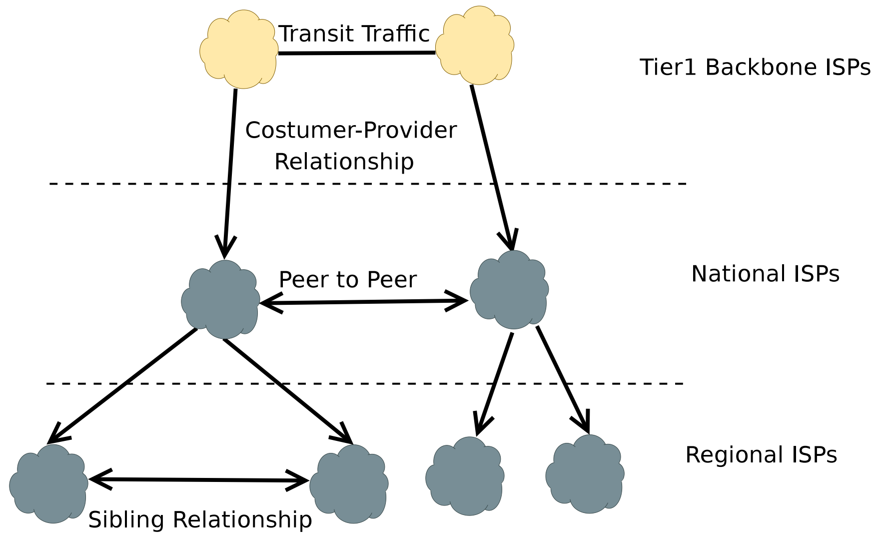

Figure 2.4: Commercial relationships between ISPs.

BGP routing update export rules that are different for the individual relationships. The proposed solution by Gao is based on forming annotated graphs of the network and making sure the AS paths are *Valley-Free*, i.e., after traversing a provider-to-customer or peer-to-peer edge (link), the AS path cannot traverse a customer-to-provider or peer-to-peer edge. The Valley-Free criteria holds only when the following conditions are met:

- A provider-to-customer edge can be followed by only provider-to-customer or sibling-to-sibling edges.
- A peer-to-peer edge can be followed by only provider-to-customer or sibling-to-sibling edges.

Subramanian et al. [SARK02] focused on peering relationships between ASes from a commercial relationship point of view. They combined BGP data from multiple vantage points to construct a view of the Internet topology, using BGP routing tables from 10 *Telnet Looking Glass* servers.[4] The proposed algorithm ranks each AS from each of the vantage points based on the number of *up-hill* and *down-hill* portions. The results suggest the design of a topology generator based on directed graphs, as opposed to degree-based methods, as the directed graphs make distinction between edge ASes, connecting to several transit core ASes.

This work led to many other interesting findings about AS-level relationships. Batista et al. [BEH$^+$07] took this approach further by proving that identifying AS relationships from BGP data, especially when measured from multiple sources, is an NP-complete problem. The suggested solution is a linear time algorithm for determining the AS relationships in the case in which the problem admits a solution without anomalies for large portions of the Internet (e.g., data obtained from single points of view). The solution is performed by starting from a set of AS paths, so that the number of invalid paths is kept small. This method can be applied on the address prefix of the hosts within an AS.

[4] http://www.traceroute.org/#LookingGlass

When looking at the path taken between ASes, direct access to end points is not always possible. The approach of using multiple sources of data is an extremely useful method in such scenarios. It enables a more detailed analysis of the possible paths between two end nodes (ASes in this case). Mao et al. [MQWZ05] explored the feasibility of inferring AS paths by using BGP tables from multiple vantage points, router-level paths from traceroute servers, and AS-level paths from Looking Glass sites.

One of the inherent issues of inference of AS-level topology of the Internet by use of mapping node IP addresses to registered AS numbers is that sibling relationships are missed. Dimitropoulos et al. [DKF+07] proposed an alternative solution to AS-level map inference which attempts to find sibling-to-sibling (s2s) relationships, as well as customer-to-provider (c2p) or provider-to-customer (p2c). The proposed inference model avoids the mistake of considering siblings as customers or peers, which in turn may result in wrong inference of a provider as a customer, or the other way round, while still rendering a path as valid. The inference of s2s links plays an important role when looking at corporate networks, where multiple ASes belong to the same organisation. In order to look at the s2s relationships, the IRR databases are consulted and dictionary of synonymous organisations is manually created. Although a disadvantage of this approach is that the IRR are not always up-to-date.

Using public BGP data and validating their results against cooperating ISPs, the author's main conclusion is that with BGP derived inference, it is possible to identify less than 50% of peer to peer links. Another conclusion is that nearly all relationships are p2p and c2p, as confirmed by the conducted survey.

When focusing on AS-level graphs of the Internet, peering relationships play an important role in providing alternative routing and resilience. Muhlbauer et al. [MFM+06] focus on the connections between the ASes within the Internet, due to the importance of the inter-AS relationships. Peering relationship are difficult to infer due to the business nature of this information, and the limited ability of methods to correctly identify such peering relationships. However their importance is significant as they affect inter-domain routing policies. They build a simple model that captures such relationships by using BGP data from observation points such as Routeviews and RIPE. They then use simulations to provide an AS-level map which they compare with the BGP data from other vantage points.

In a view inspired by the business relationships of providers, Chang et al. [CJW06] present a model of economic decisions that an ISP or AS has to make in order to peer with other ASes, or with transit tier-1 ASes. The economic decisions which have to be considered by an ISP are of three types: peering, provider and customer. In each case, the cost-centric multilateral decision, as referred to by the providers, has to bring mutual benefits for both parties. The gravity model [Poy63] has been used to describe decisions on traffic demand and exchange. The distance of ASes from each other plays the critical role in the decision made by an AS to peer with another. They use BGP data to form node degree distributions to infer peering relationships. An important result of their work is an analysis of changes in the topology of a network, by introduction of new peering relationships and updates to the current ones.

Muhlbauer et al. [MUF+07] investigated the role and limitations of business relationships as a model for routing policies. They observe that popular locations for filtering correspond to *valleys* where no path should be propagated according to inferred business relationships. This result reinforces the validity of the valley-free property used for business relationship inference. This work reveals two dimensions to policies: (i) which routes are allowed to propagate across inter-domain links (route filtering); and (ii) which routes among the most preferred ones are actually chosen (route choice) and thus observed by BGP monitors. They use BGP data from more than 1,300 BGP observation points, including Route-views. The observation points are connected to more than 700 ASes with some feeds from multiple locations. They provide a model of ASes and have identified sets of per-prefix policies in order to obtain agreement between the routes selected in their model and those observed in the BGP data.

2.4 Models of Internet Topology

Mathematical modelling of the characteristics of the Internet is a key stage for successful generation of realistic topologies. These mathematical models can range from geographical distance and clusters to distribution of nodes with different degrees of connectivity. In reality, the constant change in the Internet topology makes it difficult to obtain a single topology of the Internet and instead it is more appropriate to refer to the obtained maps as *Internet topologies*.

This section presents some of the models of the Internet topologies. The objective of this section is to familiarise the reader with the common methods of characterising the topology of a network and provide a basic understanding of the most common terms used in this context.

2.4.1 Random graphs

Complex networks such as the Internet have traditionally been described using the random graph theory of Erdös and Rényi [ER85]. In a simple model, for a given number of nodes n, edges m and the average degree $\bar{k} = 2m/n$, one can construct the class of random graphs having the same average degree k by connecting every pair of nodes with probability $p = k/n$.

Despite the ease of use of the random network model, and their ability to produce some of the required metrics for a generator such as average node degree, they were abandoned in favour of models that capture the statistical characteristics of the Internet as discussed in the next section.

2.4.2 Power laws in topologies

Power laws are one of the most widely used notions in the context of topology analysis of the Internet. Power laws are seen in statistical distributions where there is no concept of scale variance, i.e., a property, such as a distribution of nodes in a network, follows the same rules at different scales or resolutions. In a seminal paper, Faloutsos *et al.* [FFF99] stated that certain properties of the AS-level Internet topology are well described by power laws. In this work, the authors use three Internet instances (topologies inferred from BGP tables). Three specific power laws were observed and these were believed to hold for the Internet:

- *Rank exponent*: Out-degree of a node is proportional to its rank to the power of a constant.

- *Out-degree exponent*: The frequency of an out-degree is proportional to the out-degree to a constant power.

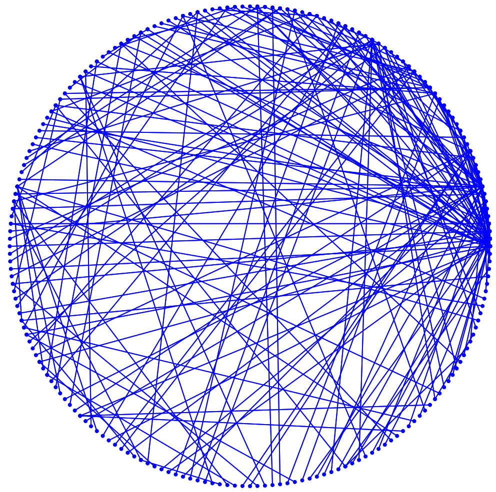

Figure 2.5: A power law network of 200 nodes

- *Eigen-exponent*: The eigenvalues of the adjacency graph are proportional to the order i to a constant power.

One of the classic models that is used in this context is the *BA* model, introduced first by Barabási and Albert [BA99]. This model is based on the incremental growth of networks, by addition of new nodes and preferential attaching nodes to well-connected ones. They also reported that Internet has power law characteristics, alongside the findings of Faloutsos et al. Barabási and Albert focus on WWW webpages and links between them as an alternative measurement of the Internet.

Figure 2.5 shows a network of 200 nodes connected based on the BA model. Such a graph will have power law characteristics, and a tree-like structure due to its scale-free nature. If one relies on the traceroute tool, it is difficult to infer the cross links between the nodes. A scale-free network is not a homogeneous network as the nodes have a very heavy-tailed distribution. Despite the small size of the Internet at the time of observations of Faloutsos *et al.*, these observations were believed to hold in future growth stages of the Internet. This hypothesis intrigued Siganos *et al.*to repeat the above analysis again [SFFF03]. They prove the existence of power laws in Internet at AS-level, looking at two topology measurements, at few snapshots over five years, one from Oregon RouteViews and another is the dataset used by Chen *et al.* [CCG$^+$02]. The test for the existence of power laws is carried on the metrics such as rank exponent, degree exponent and eigenvalues. The conclusions are that the power laws exist over a five year period and they are an efficient way to describe metrics of topology graphs.

Figure 2.6 displays the node degree distribution of the power law network in Figure 2.5, plotted on a log-log graph. Existence of a straight line indicates the existence of a power law distribution of node degrees.

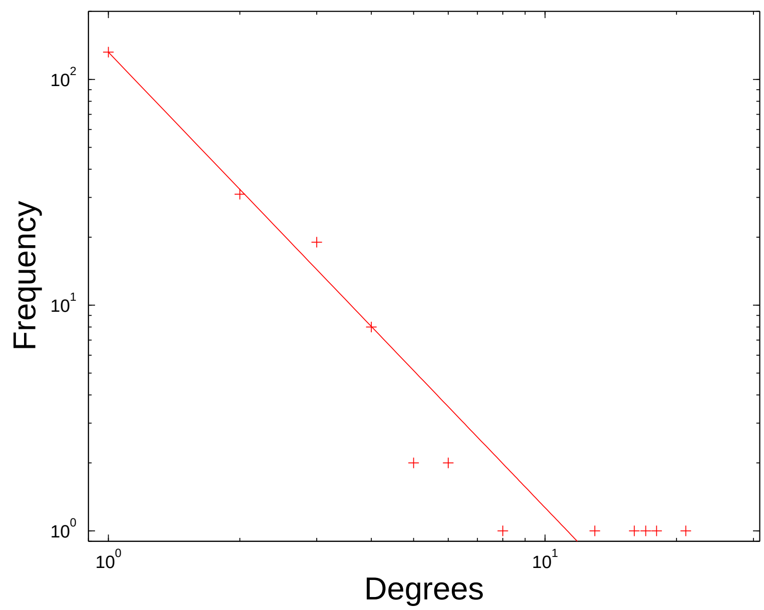

Figure 2.6: Power law node degree distribution.

The existence of power laws in the Internet is interesting as the Internet is formed from smaller networks which are self-managed. Medina *et al.* [MMB00] look at four factors in formation of Internet topologies which may cause various power laws inferred on the Internet:

1. Preferential connectivity of nodes to nodes with more connections.
2. Incremental growth of the networks.
3. Distribution of nodes in space (random or heavy-tailed).
4. Locality of edge connections (preference to connect to nearby nodes).

The BRITE topology generator[MLMB01] was used by Medina et al. to test these hypothesis. Topologies of between 500 to 15,000 nodes were considered, with and without incremental growth and preferential connectivity.

The final conclusions are that the rank and out-degree power laws are more effective in distinguishing topologies than the number of hops between nodes and eigenvalue power laws which are observed similarly in all topologies. Preferential connectivity and incremental growth are found to be the main causes for all power laws in the simulations. They establish that for best correlation coefficients (approaching 1) and slope of linear fits for rank exponents (approaching 0.5 observed by Faloutsos et al. [FFF99]) both preferential connectivity and incremental growth must be present. This methodology can be extended by grouping nodes into administrative domains.

The findings in this section indicate the existence of power laws in various statistics extracted from the Internet. However the inferred statistics are not always perfect as one cannot obtain a single snapshot of the Internet topology and must rely on various measurement techniques. I now present results which indicate that the existence of power laws are merely a side-effect of poor inference techniques.

2.4.3 Arguments against power laws

The inherent biases of traceroute sampling and collection of BGP data from limited vantage points made researchers question the true existence of power laws in the Internet AS-level topology. Chen *et al.*. [CCG+02] state that BGP data represents a partial view of the Internet, hence power laws may not exist in the strict form suggested by Faloutsos *et al.* [FFF99] for the degree distribution. This argument is based on their findings that BGP AS paths do not completely capture the topology and the data from Routeviews suggest that the node degree distribution is perhaps heavy-tailed (close to Weibull distribution) and perhaps only the tail exhibits power laws. The authors use BGP routing tables of 41 ASes and information from Looking Glass websites to infer the local AS connectivity map and compare it to the one achieved by Routeviews. Data from the European Internet routing registry (RIPE), which has the peering relationships of most European ASes, is used in order to find relationships which are not seen from BGP inference, such as siblings [CGJ+04].

Another observation in conflict with the existence of power laws is the important observation made by Mahadevan et al. [MKF+06]. For a comparative study, three distinct data source are used:

1. Traceroute data from the CAIDA Skitter project, using the 31 daily graphs for the month of March 2004.

2. Routeviews BGP data for March 2004, including static table and updates.

3. RIPE WHOIS database dump for April 07, 2004.

The findings confirm that the Skitter data displays power law characteristics [FFF99], however the WHOIS graph has an excess of medium degree nodes and hence its node degree distribution does not follow power laws. They also compared many metrics of the Skitter and RouteViews graphs to those graphs generated based on Power-law Random Graphs (PLRG) [ACL00] and it is observed that the PLRG model fails to accurately capture the important properties of the skitter or RouteViews BGP graphs. Similarly, the PLRG model fails to recreate the WHOIS graph since its node degree distribution does not follow a power law at all.

Krishnamurthy et al. [KFC+05] introduce graph sampling, in order to reduce the size of inferred topologies for analysis while preserving metrics, in this case power laws and slope of graphs. They model the network as an undirected graph at AS-level. They propose sampling the graph by deleting nodes and links probabilistically, or by contracting the graph at steps, or by generating a subset of graphs from traceroute paths. They perform probabilistic deletion of nodes and edges and can reduce the graphs by about 50-70% while keeping metrics such as power laws within an acceptable range.

2.4.4 Alternative topology models

Power laws were not the only point of interest for network researchers who used datasets from various inference projects. For example the graphs produced by Rocketfuel and Skitter consist of physical connectivity of Internet routers for an ISP or a section of the Internet. However for an improved understanding of the physical infrastructure of the Internet, it is essential to have more information about the common characteristics of links such as the link bandwidth, router capacities and etc. These concerns were first raised by Alderson et al. [ALWD05], where they focus on annotated graphs of the Internet at the IP layer with addition of bandwidth and buffer sizes. The Abilene[5] and Rocketfuel maps are used to look at various differences between network models, by use of a metric proposed as *network performance*, defined as the maximum throughput of a network under a gravity model of end user traffic demands. Hence their proposed design for designing an ISP network graph is referred to as *Heuristically Optimal Topology* which is based on having sparsely connected high speed routers at the core of the network, supported by hierarchical tree-like structure at the edges. This is similar to the proposed *Highly Optimised Tolerance* approach suggested by Carlson and Doyle et al. [CD00] and *Heuristically Optimised tradeoffs* considered by Fabrikant et al. [FKP02].

The authors propose that detailed study of the technological and economic forces shaping the router-level topology of a single ISP provides convincing evidence that the Internet is not necessarily formed of highly connected routers in the core of the network. They expect border routers again to have a few relatively high bandwidth physical connections supporting large amounts of aggregated traffic. In turn, high physical connectivity at the router-level is again expected to be confined to the network edge. They also note that modelling router-level robustness requires at a minimum adding some link redundancy (e.g., multi-homing) 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.

2.4.5 Structural models of the Internet

Alongside power laws, other metrics of network topologies have been studied extensively in the literature. One of the most important factors that has already been explained in this section is the clustering of nodes. Clustering has been widely studied using techniques of finding the clustering coefficient of the nodes in a network. An alternative to this method is *spectral filtering*. Gkantsidis et al. [GMZ03] perform a comparison of clustering coefficients, by using eigenvalues of adjacency matrices from various BGP data of networks, and also on methods of topology generation, such as BRITE. This work identifies a global problem with topology generators; inability to generate representative topologies. Use of a small topology leads to concentrating only on degree distribution power laws in AS and router-level geographic topologies, as opposed to looking into the peering relationships, clustering and amount of traffic on the links. They have introduced the basics of degree-based graph generation and conditions that the links and nodes are attached to ensure connectivity, using a Markov-chain-based algorithm.

They believe that degree-sequence is not sufficient for topology generation that matches the real data. They use clustering methods and eigenvalues to analyse the generated topologies and compare with real data from NLANR[6]. The generation methods that meet a degree-sequence while incorporating clustering are suggested by the researchers. Good clustering methods are also needed in topology generators, as both the degree-sequence and the clustering are found in real networks [GMZ03].

Li et al. [LAWD04] discuss the need for topology inference and generation at different levels. For

[5] `http://abilene.internet2.edu/`
[6] `http://www.nlanr.net/`

congestion control protocols, IP level connectivity with bandwidth and buffer sizes is needed, while for attack assessment and network planning a detailed map of node and router capacities is required. For routing protocols one needs a graph of AS-level connectivity and peering information. The authors focus on node degree distribution and their heavy-tailed characteristics and whether the node degree distribution is the most important objective of a topology. They discourage the use of random generators as they do not produce power laws in node degrees, so they have been replaced by degree-based generators. The proposed *first principles approach* focuses more on physical layer, router and links. In the context of network engineering for an ISP, physical metrics such as performance and likelihood are used to for graph generations. They observe that simple heuristically designed and optimised models that reconcile the tradeoffs between link costs, router constraints, and user traffic demand, result in configurations that have high performance and efficiency.

The Internet has a hierarchical structure in the form of different tiers. Jaiswal et al. [JRT04] look at comparing the structure of power-law graph generators and that of the Internet AS graph. This is an important step in proving the existence of power laws. By decomposing graphs of the Internet at different levels, the authors establish the properties of power-law graphs and the Internet graph and find skewed distributions in degree connectivity, i.e., a large number of less-connected nodes connect to the well-connected ones, and well-connected ones tend to interconnect more closely.

Carmi *et al.* [CHK$^+$06] use the data from the DIMES project, combined with AS-level maps from the RouteViews project, to form a map of the Internet. The map formation method is based on k-*shell decomposition*, which involves removing nodes in groups based on number of connections they have, to form shells of nodes. In the first step, the k-*pruning* technique is performed by removing all the nodes with only one neighbour recursively, as well as removing the link to that neighbour along with the node. The nodes removed in this step are called the 1-shell. This process carries on with index k to form shells of higher connectivity degree. The last nonempty k-core will be, by definition, the backbone of a network such as Internet. Figure 2.7 displays a sketch of the k-core decomposition for a small graph from Alvarez-Hamelin *et al.* [AHDBV06]. Each closed line contains the set of vertices belonging to a given k-core, while colours on the vertices distinguish different k-shells.

Carmi *et al.*found that for the DIMES data used, the size of each k-shell decreases with a power law distribution, $n(k) \propto k^{-\delta}$, where the exponent δ is about 2.7.

Node clustering techniques have also been used to characterise Internet topologies. Wool and Sagie [SW04] propose a clustering method that enables the view of Internet topology as AS-graphs in different granularity levels. They find few main dense cores, which inter-connect the regional cores. They compare various degree-based generators and state the need to consider power laws and clustering coefficients when generating topologies in BRITE and Inet. They use the dense k-subgraph approach for clustering in different levels.

Yook et al. at [YJB02] propose a model of networks based on fractals. They find that the physical layout of nodes form a fractal set, determined by population density patterns around the globe. The placement of links is driven by competition between two models: preferential attachment and linear

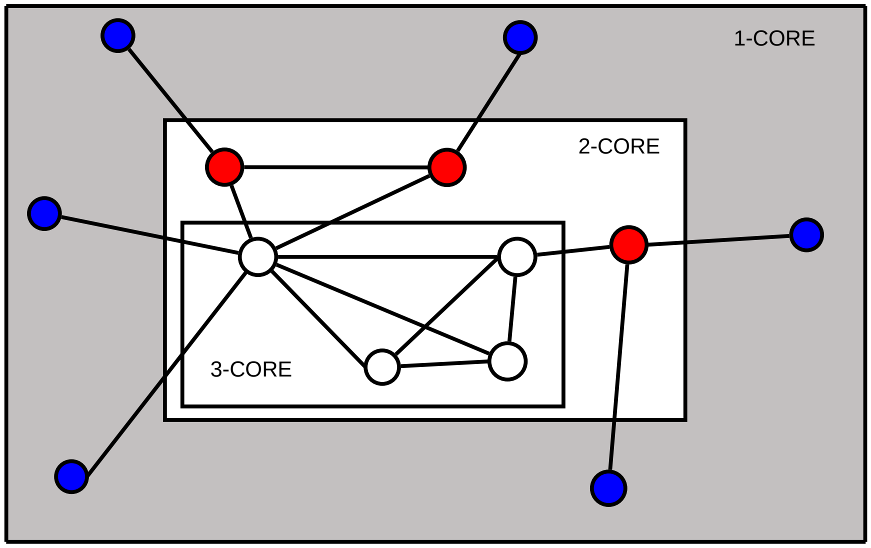

Figure 2.7: Example of k-core decomposition [AHDBV06]

distance dependence. Preferential attachment assumes that the probability that a new node will link to an existing node with k links depends linearly on k. The nodes with higher connectivity degree are more desirable for attachment by new nodes. Preferential attachment is believed to be one of the main reasons for power-law properties of the Internet. Linear distance dependence is due to the fact that the further the nodes are from each other, the less likely it is for them to have a direct connection.

The Internet, like many complex networks, is believed to have small world characteristics. Such characteristics are important for delivery of messages and content on networks. Jin and Bestavros [JB06] consider the small world characteristics when generating topologies at router-level and AS-level. At AS-level, the high variability in node degree, and at router-level the preference for local connectivity results in this phenomena. They use simulation of multicast trees on different models. They also use AS graphs of the University of Michigan AS graph dataset (RouteViews plus Looking Glass), and various router-level graphs including Skitter. They use these to get the statistics such as node degree and local connectivity in order to evaluate their model. They suggest simulators taking into consideration vertex degree distributions as well as preference for local connectivity and suggest improvement by considering scale-free characteristics as well.

The Internet architecture and structure is constantly evolving. Pastor-Satoras and Vespignani [PSV04] highlight the self-organising nature of the Internet and its evolution since birth from a statistical and physical view point. Their conclusion is that the Internet can be modelled as a network of nodes and links growing in a scale-free manner. However the growth and death rates of ISPs and ASes and predictions for future trends on the Internet remain open issues.

This section has gathered various models that are presented for the Internet at physical and routing levels. The variety of models is an indication of the complex structure of the Internet which makes it difficult to capture all the characteristics with a simple model. Based on these models, researchers develop topology generators which are discussed in Section 2.5.

2.4.6 Comparison of topology generation models

Despite the availability of many topology models, there has not yet been an agreement between researchers on a single standard method of modelling and generation of Internet or ISP network topologies. This inconsistency is due to the many aspects that one has to consider when studying a topology. In addition, different models may be used by researchers depending on the level of complexity required.

Chang et al. [CJW03] look at the problem of generating AS-level topology of the Internet. They discuss the weakness of current power-law based generators and BGP-inferred AS topologies in detecting AS peering and business relationships. The authors focus on the optimisation of a topology based on AS geography, business model and evolution in time, using the RouteViews data plus inferred information from Looking Glass sites to form two datasets. For simplicity, all multi-homing and multiple connections of ASes are removed by choosing just one link based on criteria such as lowest average hop distance. The final graph is one which is 50% of the size of original dataset, with similar node degree distribution.

Alderson et al. [ADGW03] discuss generating topologies using the Highly Optimised Tolerance concept. In this strategy, the focus of the generator is the economic trade-offs, such as cost and performance, and technical barriers faced by an ISP when designing its own network. This would allow for a focus into economical challenges faced by network operators. These issues are important for backbone service providers, which must ensure optimised use of the network capacity.

Mahadevan et al. [MKFV06] discuss the lack of analysis and topology generation tools that can focus on specific requirements of metrics of a graph, focusing on degree correlations of subgraphs of a graph that represents a network or Internet. However this method becomes extremely complex as the number of correlated nodes increases. In a basic model, a set of subgraphs are defined with various distributions and are used to define a topology. The metrics considered for analysis are: spectrum, distance distribution, betweenness, node degree distribution, likelihood (sum of products of degrees of adjacent nodes) and clustering. However in practice, the focus has been put on connectivity as the other metrics are hard to compare and classify. They focus on reproducing a given network topology and compare their results with the Skitter dataset and BGP data from RouteViews.

Mahadevan et al. believe an improvement in topology generation can be achieved by focusing on peering relationships and graph annotations such as bandwidth, latency and etc. In Orbis [MHK+07], the aim is to produce a random graph of desired size while keeping the characteristics of the input graph. They allow a user to feed in average degree, node degree and joint degree distributions from a measured topology, and the tool should also annotate the routers with AS memberships and annotate the AS links with type of relationship between them.

They observe that the AS-level topologies can be approximated by power laws. However the router-level topology does not fit strict power laws. The observed maximum degree at router-level does not increase significantly by increasing the size of the graph. In $1k$-rescaling, they attempt to preserve the shape of the PDF of the graph's degree distribution. In $2k$-rescaling, they try to preserve the degree correlation profile. They encourage the addition of latency and bandwidth distribution as another metric for rescaling for realistic topology generation.

One of the objectives of generation of topologies which closely map those of Internet is to arm network researchers with tools with which they can analyse various issues in and around the Internet, such as congestion, optimal routing and fault finding. Spring et al [SMA03] look at traceroute measurements, using *scriptroute*, from around 40 vantage points on planetlab to look at topology and routing policies internal and between ISPs to analyse the causes of path inflation, and find that inter-domain routing and peering policies have significant effect on the inflation. They suggest improvements to BGP policies to look after routing across ISPs, as the ISPs have to use minimum AS hop-count which may take longer sometimes. They compare the taken routes to the topology that they inferred using Rocketfuel.

2.5 Topology Generation

For successful simulations of traffic and network events, any generated network model must be topology aware. Topology generation is an area which researchers have been actively working on in the last decades. The first generated topologies were randomly generated by selecting a certain number of nodes and randomly assigning links between them. This was due to the lack of understanding of the architecture of the Internet and the lack of validation tools. In this section, some of the popular network topology generators are discussed.

2.5.1 Waxman

The Waxman model of random graphs is based on a probability model for interconnecting nodes of the topology given by:

$$P(u,v) = \alpha e^{-d/(\beta L)} \tag{2.1}$$

where $0 < \alpha, \beta \leq 1$, d is the Euclidean distance between two nodes u and v, and L is the network diameter, i.e., the largest distance between two nodes. Note that d and L are not parameters for the Waxman model. The Internet is known not to be a random network but I include the Waxman model as a baseline for comparison purposes. Figure 2.8 displays a topology generated by the Waxman model. It can be seen that some nodes are not connected to others.

2.5.2 GT-ITM

With the explosion of the Internet, researchers realized that they need to capture the structural properties and attempted to model the design of the Internet. The hierarchical modelling of the Internet topology was originally done by the transit-stub models. Calvert et al. [CDZ97] presented one of the first results in this field by focusing on the graph-based models to represent the topology. The parameters used include the number of transit and stub domains, number of Local Area Networks (LANs) per stub domain, and the number of edges (links) between transit and stub domains, to initialise the topology generator. Then the transit domains, transit nodes and their inter-connecting edges are placed and similarly the stub domains. The Transit-Stub model produces connected subgraphs by repeatedly generating a graph according to the edge count and checking the graph for connectivity. Unconnected graphs are discarded.

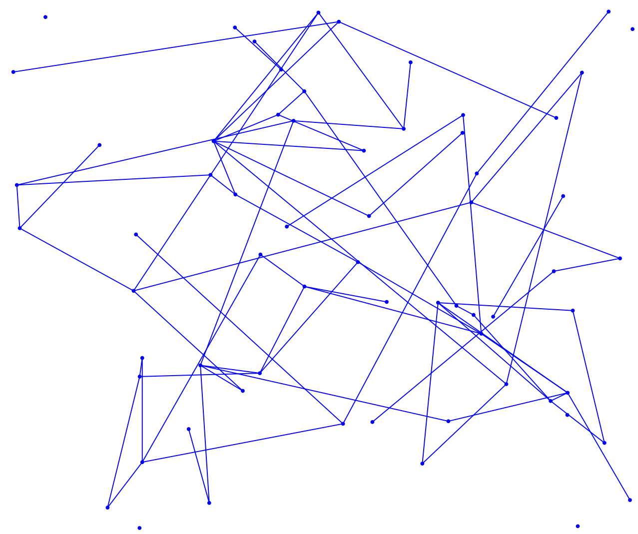

Figure 2.8: A topology generated by the Waxman model.

This method ensures that the resulting subgraph is taken at random from all possible (connected) graphs; however, it may take a long time to generate a connected graph if the edge count is relatively small compared to the number of nodes. Extra edges from stub domains to transit nodes are added by random selection of the domains and nodes.

Georgia Tech Internetwork Topology Models (GT-ITM), also known as the *Transit-Stub* generator, is capable of producing several forms of network topologies:

- *Flat random graphs*: GT-ITM has five models of topology embedded within it including pure random model and varieties of the Waxman [Wax88] model. These are not hierarchical models.

- *N-Level model*: The N-Level model constructs a topology recursively. In this method, a graph is made by dividing the Euclidean plane into equal-sized square sectors, and then each sector is divided into smaller sectors in the same manner, so the scale of the final graph is equivalent to that of the individual levels.

- *Transit-Stub model*: This model produces interconnected transit and stub domains. This model is controlled by number of domains, average node per transit domain, average stub domains per transit domain, and average nodes per stub domain.

In the transit-stub domain, care has been taken to ensure that the paths are similar to those of the Internet, for example the path between two stub domain goes through one or more transit domains and not any stub domains and not the other way round. This is done by assigning appropriate weights to the interdomain edges.

The transit-stub model is comparable to the *Tiers* model [Doa96], in which the three levels of hierarchy, or tiers, are referred to as Wide Area Network (WAN), Metropolitan Area Network (MAN), and

LAN levels, corresponding to the transit domains, stub domains, and LANs of the transit-stub method. The Tiers model produces connected subgraphs by joining all the nodes in a single domain using a minimum spanning tree algorithm, a popular method used as the basis for laying out large networks. This generation method has been tried in two implementations of Transit-Stub (TS) model, part of GT-ITM.

2.5.3 BA and AB

These models are inspired by the Barabasi and Albert [BA99] model of networks, and the Albert and Barabasi (AB) model based one evolving networks [AB00] which incorporate preferential attachment and incremental growth factors. Starting with a network of m_0 isolated nodes, $m \leq m_0$ new links are added with probability p. One end of each link is attached to a random node, while the other end is attached to a node selected by preferring the more popular, i.e., well-connected, nodes with probability

$$\Pi(k_i) = \frac{k_i + 1}{\sum_j k_j + 1} \tag{2.2}$$

where k_j is the degree of node j, with probability q, m links are rewired and new nodes are added with probability $1 - p - q$. A new node m has m new links that, with probability $\Pi(k_i)$, are connected to nodes i already present in the system.

2.5.4 GLP

The Generalised Linear Preference model (GLP) [BT02] focuses on matching characteristic path length and clustering coefficients. It uses a probabilistic method for adding nodes and links recursively while preserving selected power law properties. In the GLP model, when starting with m_0 links, the probability of adding new links is defined as p where $p \in [0, 1]$. Let $\Pi(d_i)$ be the probability of choosing node i. For each end of each link, node i is chosen with probability $\Pi(d_i)$ defined as:

$$\Pi(d_i) = (d_i - \beta)/\sum_j (d_j - \beta) \tag{2.3}$$

where $\beta \in (-\infty, 1)$ is a tunable parameter indicating the preference of nodes to connect to existing popular nodes.

2.5.5 Inet

Inet produces random networks using a preferential linear weight for the connection probability of nodes after modelling the core of the generated topology as a full mesh network. Inet sets the minimum number of nodes at 3037, the number of ASes on the Internet at the time of Inet's development. By default, the fraction of degree 1 nodes α is set to 0.3, based on measurements from Routeviews[7] and NLANR[8] BGP table data in 2002.

2.5.6 The Positive Feedback Preference (PFP)

In the Positive Feedback Preference (PFP) model [Zho06], the AS topology of the Internet is considered to grow by interactive probabilistic addition of new nodes and links. The PFP model starts with a random network of size n. At each time step:

[7] `http://www.routeviews.org/`
[8] `http://www.nlanr.net/`

1. With probability p, a new node is attached to a host node, and at the same time a new internal link appears between the host node and a peer node.

2. With probability $q \in [0, 1-p]$, a new node is attached to a host node, and at the same time two new internal links appear between the host node and two peer nodes.

3. With probability $1-p-q$, a new node is attached to two host nodes, and at the same time a new internal link appears between one of the host nodes and a peer node.

2.5.7 IGen

Another generator which aims to generate topologies which have the geographical problems associated with network design is the Igen generator. Quoitin [Quo05] explains why it is difficult to infer topologies and thus proposes the generation of topologies based on network design parameters. He argues why pure degree-based generators such as BRITE or GT-ITM fail to capture real optimisation challenges faced by network designers. The metrics such as latency minimisation, dimensioning and redundancy are discussed. IGen first creates PoPs to look like the Sprint network, then it make connected trees based on the Highly Optimised Tolerance methodology [ADGW03].

2.6 Summary

Internet's complex architecture and organisational structure hinders the construction of accurate maps of the network and makes it nearly impossible to propose definitive mathematical models. Understanding the network at the physical layer is essential for routing and resilience purposes. understanding the higher layers, the virtual types of connectivity structures are very different when studied from different sources of data and a correct understanding of the nature of these connections is essential for traffic engineering and economic modelling of the network.

The research efforts towards mapping the internet have focused on trying to get a map at router and AS level. Researchers try to understand routing policies and provide connectivity maps, by focusing on the router and AS-level graphs.

The development of the above works suggest that realistic topology generators will benefit from taking link bandwidth and geographic distribution of the nodes into consideration. It is also becoming increasingly important for network researchers to take into consideration the evolution and structure of networks and Internet as a whole over time and the presence of annotated links plays an important role in this context.

In this chapter I have briefly introduced the challenges in different areas of Internet topology research. In Chapter 3, I put the available AS topology models under test, comparing them at different network sizes with observed Internet topologies. In Chapter 4, I introduce a new metric for tuning the parameters of topology models in order to be comparable to observed datasets from different measurement infrastructures. I also study the evolution of the Internet in Chapter 5, analysing the effects of measurement biases on our understanding of the Internet topology.

Chapter 3

Understanding Internet AS Topology Models

Many models have been proposed for generating Internet Autonomous System (AS) topologies, most of which make structural assumptions about the AS graph. In this chapter I compare topologies generated from several different models against a set of measured AS topologies. In contrast to past work, I avoid making assumptions about which topological properties are important for characterising the AS topology by using a large set of topological metrics in the analysis.

In this chapter I show that current topology generators fail to capture the complexity of the local interconnection structure between ASes, despite matching degree-based properties of the AS topology reasonably well. Using a collection of BGP topologies from many measurement locations, I also analyse the reference datasets. I observe that adding more measurement locations significantly affects, especially in the core, local structure properties such as clustering and node centrality while not notably affecting degree-related metrics. The failure of topology generators thus stems from an underestimation of the importance of the complexity of connectivity in the core caused by inappropriate use of BGP data.

3.1 Introduction

For many years, researchers have modeled the Internet's Autonomous System (AS) topology[1] using graphs obtained via various measurement techniques such as BGP routing tables [Hal97, RLH06b] and traceroute maps [HPM$^+$02]. The AS topology is an abstraction of the Internet commonly used to analyse its characteristics such as size and connectivity patterns, and to simulate the effects and performance of new protocols.

Figure 3.1 illustrates the relationship between the real Internet topology, measurements of it and the topology generation models which are discussed in this chapter. Observations of the AS topology suffer from two problems: a given set of observation points has only limited visibility of the topology, and each observation technique suffers from measurement artifacts. In this chapter I treat observations from BGP and traceroute as samples of reality, accepting that they suffer from biases and reveal different partial truths about the properties of the Internet.

At the same time, the models which underlie topology generators make simplifying assumptions about the topology [BT02, MKFV06, Zho06] based on prior observations. At present, the main widely-

[1] Note that the AS topology does not represent the data-plane topology. Many organisations are permanently connected to their providers, sharing an AS number [SBCC98]. Alternately, an organisation may use many AS numbers for controlling routing.

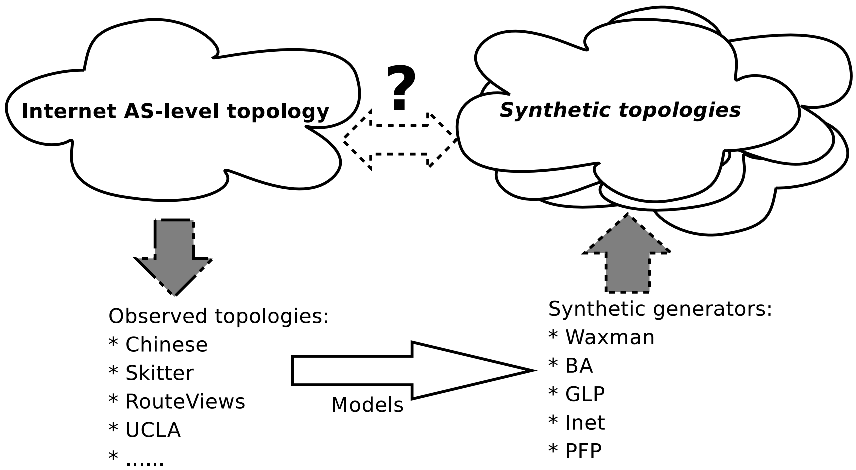

Figure 3.1: Internet topology generation

held assumptions are that the AS topology has a hierarchical structure and its node-degree distribution obeys a power-law. Note that correct reproduction of the hierarchical structure can be achieved simply by following degree-related distributions [TGJ+02], although both the node degree distribution and the joint degree distribution must be reproduced [MKF+06]. Thus, by comparing different observed topologies with different levels of incompleteness, with topologies generated from different models, I learn about the limitations of particular assumptions about the Internet's AS topology. The direction of these biases and limitations gives us insight into the actual properties of the AS topology.

In this chapter I show that current topology generators capture the node degree distributions quite well, but fail to account either for the complex local interconnection structure between ASes, or the highly meshed structure of the core AS topology. Such failures can affect the performance of protocols and applications when simulated using synthetic topologies. For example a routing protocol can demonstrate different convergence times on a random graph when compared to a graph with high number of alternative paths between nodes.

Different metrics are considered important by different topology generatorion models, so comparing topologies from different generators requires taking a broad perspective. A key principle underlying this work is to be agnostic about the topological properties of the Internet. To this end I use many topological metrics, hoping to cover a large enough set of properties of the true AS topology to reveal as many biases in observations as possible. I do not claim that the set of metrics used captures all important aspects of the AS topology. However, using such an extensive set of topological metrics allows one to observe even subtle differences between synthetic topologies and observed ones. Also, I use statistical measures for comparing distributions of some metrics, allowing us to objectively compare the similarity of two topologies.

The primary purpose of topology generators is to provide *realistic* topologies for simulation, where this means that their properties should be as close as possible to those of the real AS topology. This is typically tested by comparison with measured topologies, which suffer from the biases discussed above. A further problem is that the true topology is itself dynamic: it changes due to routing dynamics, misconfiguration [FB05], and the long-term evolution of the Internet. This results in problems for BGP-based observations, as well as traceroute-based observations. For example, traceroute can report AS paths hops that do not map to a unique AS number [MRWK03][2]. Thus, alongside comparison among generators based on different underlying assumptions, I contrast the results with measurements made at different times using different techniques and observation points.

In summary, the measurements suggest that using additional BGP peers for collecting connectivity information greatly affect important characteristics such as power laws and measures of centrality, while having little affect on basic degree-related properties. This suggests that to understand the nature of the Internet topology, one should only use rich datasets which capture a large portion of peering links.

The key contributions of this chapter are to characterise the existing generators across a large set of metrics, and to compare them to numerous available measured datasets. I show that *power laws* are not strictly adhered to in today's Internet AS topology. My results also indicate that the AS topology is best modeled by matching node degree distributions while taking into consideration the meshed core formed by the many peering links between ASes. I also give insight into the effect of varying the number of observation points for capturing the AS topology.

The rest of this chapter is structured as follows. In Section 3.2, I contrast past work with my analysis. I revisit current AS topology models and describe their underlying assumptions in Section 3.3, present a collection of observed AS topologies collected using different methodologies from various locations in the world in Section 3.4. In Section 3.5, I describe commonly used metrics for topology characterisation. In Section 3.6 I discuss the appropriate statistical measures of similarity and then in Section 3.7, I present the results of the analysis. I discover that synthetic topologies and observed topologies record biases due to the nature of the data collection processes. Hence I conduct an intensive analysis of the topology dataset collected from a large number of measurement locations and analyse the impact of increasing the number of BGP peering vantage points. Alongside a description of my methodology, the results in Section 3.8 show that the importance of preferential attachment has weakened while peering links, underestimated in the past, are now far more critical. As well as concluding, Section 3.9 discusses potential improvements in the field.

3.2 Related Work

Zegura *et al.* [ZCD97] analyse topologies of 100 nodes generated using pure random, Waxman [Wax88], exponential and several locality based models of topology such as Transit-Stub. They use metrics such as average node degree, network diameter, number of paths between nodes. They find that pure random graphs produce topologies that represent expected properties such as locality very poorly and so I exclude pure random graphs from the comparisons. They suggest that the Transit-Stub method should be used due to both its efficiency and the realistic average node degree its topologies achieve.

Faloutsos *et al.* [FFF99] state that three specific properties of the AS-level Internet topology are well described by power laws: rank exponent, out-degree exponent and eigen exponent (graph eigenvalues).

[2] This effect is also seen in CAIDA's Skitter dataset, where a number of possible AS numbers are recorded for a router on the traceroute path.

This work parallelled development of many models based on power laws, such as the Barabási and Albert [BA99] model, based on incremental growth by addition of new nodes and preferential attachment of new nodes to existing well-connected nodes.

Later, Bu and Towsley [BT02] used the empirical complementary distribution (ECD) rather than standard histograms to generate new nodes. They showed the variability in graphs from different generators using the same heuristics using characteristic path length and clustering coefficients.

Tangmunarunkit *et al.* [TGJ+02] provide a first point of comparison of the underlying characteristics of degree-based models against structural models. A major conclusion is that the degree-based model, in its simplest form, performs better than random or structural models at representing all the studied parameters. They compare three categories of model generators: the Waxman model of random graphs, the Tiers [Doa96] and Transit-Stub structural models, and the simplest degree based generator, called the power-law random graph (PLRG) [ACL00]. They compare under three metrics: expansion, resilience and distortion. It was found that the PLRG performs better than the random or structural models in reproducing these parameters. Based on their defined metrics, they conclude that the hierarchy present in the measured networks is stricter than in degree-based generators. However, they leave many questions unanswered about the accuracy of degree-based generators and the choice of metrics.

Zhou and Mondragon [Zho06] propose models based on several mathematical features, such as rich-club, interactive growth and betweenness centrality. They use AS data from the CAIDA Skitter project to examine the JDD and rich-club connectivity. They show that for these data, rich-club connectivity and the JDD are closely linked for a network with a given degree distribution.

In this chapter, I consider many more recent degree-based generators using a larger set of graph-theoric metrics to give better insight into correct understanding of the AS topology. I make a detailed comparison with a range of different Internet AS topologies at national and international level obtained from traceroute and BGP data. When choosing the metrics, I considered both metrics used by the topology generator designers and those used more widely in graph theory. A particular point to note is that I chose not to use the three metrics of Tangmunarunkit *et al.* for two reasons. First, computation of both resilience and distortion are NP-complete, requiring use of heuristics. In contrast, all the metrics used in this chapter are straightforward to compute directly. Second, although accurate reproduction of degree-based metrics is well-supported by current topology generators, my hypothesis was that local interconnectivity was poorly supported, and so I chose to use several metrics that focus on exactly this, e.g., assortativity, clustering, and centrality.

3.3 AS Topology Models

There are many models available that claim to describe the Internet AS topology. Several of these are embodied in tools built by the community for generating simulated topologies. In this section I describe the particular models whose output is compared in this section. The first are produced from the Waxman model [Wax88], derived from the Erdös-Rényi random graphs [ER85], where the probability of two nodes being connected is proportional to the Euclidean distance between them. The second come from the Barabasi and Albert (BA) [BA99] model, following measurements of various power laws in degree distributions and rank exponents by Faloutsos *et al.* [FFF99]. These incorporate common beliefs about preferential attachment and incremental growth. The third are from the Generalised Linear Preference model [BT02] which additionally model clustering coefficients. Finally, Inet [WJ02] and PFP [Zho06] focus on alternative characteristics of AS topology: the existence of a meshed core, and the phenomenon of preferential attachment respectively. Each model focuses only on particular metrics and parameters, and has only been compared with selected AS topology observations.

3.4 AS Topology Observations

The Internet AS topology can be inferred from various sources of data such as BGP routing or traceroute [Mal93] at the network (IP) layer. Using just BGP routing data suffers from incompleteness, no matter how many vantage points are used to collect observations. In particular, even if BGP updates are collected from multiple vantage points and combined, many peering and sibling relationships are not observed [FMM$^+$04]. Conversely, traceroute data misses alternative paths since routers may have multiple interfaces which are not easily identified, and multi-hop paths may also be hidden by traffic tunnelled via Multi-Protocol Label Switching (MPLS) pathways. Combining these data sources does not solve all problems since mapping traceroute data to AS numbers is not always accurate [MRWK03]. In this chapter I attempt to avoid these problems by comparing against many measurement-derived datasets giving a diverse spatial and temporal comparison across different continents and years of measurement.

3.4.1 Chinese AS topology

The first dataset is a traceroute measurement of the Chinese AS Topology collected from servers within China in May 2005. It reports 84 ASes, representing a small subgraph of the Internet. Zhou *et al.* [ZZZ07] maintain that the Chinese AS graph presents all the major topology characteristics of the global AS graph. The presence of this dataset enables us to compare the AS topology models at smaller scales. Further, this dataset is believed to be nearly complete, i.e., it contains very little measurement bias and accurately represents the true AS topology for that region of the Internet.

3.4.2 Skitter

The second dataset comes from the CAIDA Skitter project[3]. CAIDA computes the adjacency matrix of the AS topology from the daily Skitter measurements. These are obtained by running traceroutes over a large range of IP addresses and mapping the prefixes to AS numbers using RouteViews BGP data. Since the Skitter data represents paths that have actually been traversed by packets to their destinations, rather than paths calculated and propagated by BGP system, it is more likely to faithfully represent the IP topology than the BGP data alone. For this study, I used the graphs for March 2004 as used by Mahadevan *et al.* [MKF$^+$06]. This dataset reports 9,204 unique ASes across the Internet.

3.4.3 RouteViews

The third dataset I use is derived from the RouteViews BGP data. This is collected both as static snapshots of the BGP routing tables and dynamic BGP data in the form of BGP message dumps (updates

[3] `http://www.caida.org/tools/measurement/Skitter/`

and withdrawals). I have used the topologies provided by Mahadevan *et al.* [MKF+06] from two types of BGP data from March 2004: one from the static BGP tables and one from the BGP updates. In both cases, they filter AS-sets and private ASes and merge the 31 daily graphs into one. This dataset reports $17,446$ unique ASes across 43 vantage points in the Internet.

3.4.4 UCLA

The fourth dataset comes from the Internet topology collection[4] maintained by Oliveira *et al.* [OZZ07]. These topologies are updated daily using the data sources such as BGP routing tables and updates from RouteViews, RIPE[5], Abilene[6] and LookingGlass servers. Each node and link is annotated with the times it was first and last observed. I use a snapshot of this dataset from November 2007 computed using a time window on the last-seen timestamps to discard ASes which have not been seen for more than 6 months. The resulting dataset reports 28,899 unique ASes.

3.5 Topology Characterisation

In this section I provide a large set of topological metrics. Taken individually, those metrics do not define a *distance* in graph space, i.e. how two graphs look like each other. However, once combined, they can identify the failures of topology models and highlight the potentials for improvements. The topological metrics are computed for the synthetic and measured topologies, modeled as graphs with a collection of nodes and undirected links that connect pairs of nodes, $G = (\mathcal{N}, \mathcal{L})$ with $N = |\mathcal{N}|$ nodes and $M = |\mathcal{L}|$ links. In the remainder of this thesis, I consider the networks formed by the largest connected component. Consequently, the computation of the topological metrics is restricted to those largest connected components of the inferred topologies.

3.5.1 Degree

The degree k of a node is the number of links adjacent to it. The *average node degree* $\bar{k}$ is defined as $\bar{k} = 2M/N$. The *node degree distribution* $P(k)$ is the probability that a randomly selected node has a given degree k and is defined as $P(k) = n(k)/N$, where $n(k)$ is the number of nodes of degree k. The *joint degree distribution* (JDD) $P(k, k')$ is the probability that a randomly selected pair of nodes has degrees k and k'. A summary measure of the joint degree distribution is the average neighbour degree of nodes with a given degree k, $k_{nn}(k) = \sum_{k'=1}^{k_{max}} k' P(k'|k)$. The maximum possible $k_{nn}(k)$ value is $N - 1$ for a maximally connected network, i.e. a complete graph. Hence, I represent JDD by normalised values $k_{nn}(k)/(N-1)$ [MKF+06].

3.5.2 Assortativity

Assortativity is a measure of the likelihood of connection of nodes of similar degrees [New02]. This is usually expressed by means of the assortativity coefficient r: assortative networks have $r > 0$ (dis-assortative have $r < 0$ respectively) and tend to have nodes that are connected to nodes with similar (dissimilar respectively) degree.

---

[4] `http://irl.cs.ucla.edu/topology/`
[5] `http://www.ripe.net/db/irr.html/`
[6] `http://abilene.internet2.edu/`

3.5.3 Clustering

Local clustering $C(k)$ is the ratio of $\bar{m}_{nn}(k)$, the average number of links over all the connected components between the neighbours of k-degree nodes, and the maximum possible number of such links $C(k) = 2\bar{m}_{nn}(k)/(k(k-1))$. I use distribution of clustering coefficients C, which is the proportion of triangles (nodes with two connected neighbours) among all connected node triplets in the entire network which gives the same weight to each triangle in the network irrespective of degree of the nodes.

3.5.4 Rich-Club

The rich club coefficient $\phi(\rho/n)$ is the ratio of the number of links in the component induced by the ρ largest-degree nodes to the maximum possible links $\rho(\rho-1)/2$ where $\rho = 1...n$ are the first ρ nodes ordered by their non-increasing degrees in a network of size n nodes [CFSV06].

3.5.5 Shortest path length distribution

The shortest path length distribution $S(h)$, as commonly computed using Dijkstra's algorithm, is the probability distribution of two nodes being at minimum distance h hops from each other. From the shortest path length distribution, the average node distance in a connected network is derived as $\bar{H} = \sum_{h=1}^{h_{\max}} hS(h)$, where $h_{\max}$ is the the shortest paths between any pair of nodes with the greatest number of hops. $h_{\max}$ is also referred to as the diameter of a network.

3.5.6 Centrality measures

Betweenness centrality is a measure of the number of shortest paths passing through a node or link, a centrality measure of a node or link within a network. The betweenness for a node is $B(v) = \sum_{s \neq v \neq t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}}$ where σ_{st} is the number of shortest paths from s to t and $\sigma_{st}(v)$ is the number of shortest paths from s to t that pass through a node v [HKYH02]. The average node betweenness $\bar{B}$ is the average value of the node betweenness over all nodes $\bar{B} = \sum_{v=1}^{n} B(v)$.

Closeness is another measure of the centrality of a node within a network and is defined as the average length of the shortest paths to and from all the other nodes in a graph. The closeness $S(v)$ for a node v is the reciprocal of the sum of shortest paths to all other reachable nodes (connected component) V in a network $S(v) = \frac{1}{\sum_{v,t \in \mathcal{N}} V}$. A high closeness of a node is indicative of it having short geodesic distance to other nodes [Sab66].

3.5.7 Coreness

The l-core of a network is the maximal component in which each node has at least degree l. In other words, the l-core is defined as the component of a network obtained by recursively removing all nodes of degree less than l. A node has coreness l if it belongs to the l-core but not to the $(l+1)$-core. Hence, the l-core layer is the collection of all nodes having coreness l. The core of a network is the l-core such that the $(l+1)$-core is empty [BBGW04].

3.5.8 Top clique size

A clique in a network is a set of pairwise adjacent nodes, i.e., a component which is a complete graph. The top clique size, also known as the graph clique number, is the number of nodes in the largest clique in a network [Woo97].

3.5.9 Spectrum

The spectrum is the set of eigenvalues of the adjacency matrix of a graph. Recently, it has been observed that eigenvalues are closely related to almost all critical network characteristics [Chu97]. For example, Tangmunarunkit *et al.* [TGJ$^+$02] classified network resilience as a measure of network robustness subject to link failures, resulting in a minimum balanced cut size of a network. Spectral graph theory enables studying network partitioning problem using eigenvalues [Chu97].

In the graph theory literature, one usually considers the adjacency or the Laplacian matrix [Mer95, CDGT88], which employ different normalisation and therefore lead to different spectra. In this chapter I focus on the spectrum of the *normalised Laplacian matrix* [Chu97], where all eigenvalues lie between 0 and 2, allowing easy comparison of networks of different sizes.

3.6 Measures of Similarity

To compare the distributions of various metrics I use the following statistics to determine how close two distributions are to each other. I perform the calculations for each synthetic topology instance separately and compare them to observed topologies of the same size. Note that distances are relative to the metric and the topology size, and so the distances of one metric for a particular sized topology cannot be compared either to distances of another metric for the same sized topology, or to distances for the same metric for different sized topologies.

3.6.1 Kolmogorov-Smirnov (KS) distance

Given samples of two random variables, X_1 and X_2, the KS distance is the maximum empirical distribution difference defined as:

$$D_{max} = \text{sup}\ |F_{n_1}(x) - F_{n_2}(x)|$$

where $sup\ S$ is the supremum of set S and $F_{n_i}(x)$ is the empirical distribution of $X_i (i = 1, 2)$:

$$F_{n_i}(x) = \frac{1}{n_i} \sum_{j=1}^{n_i} I_{X_j \leq x} \text{ for } i\text{=1,2}$$

where n_1 and n_2 are the number of samples from X_1 and X_2 and I_{X_j} is the indicator function.

The closely related 2-sample KS test tests the null hypothesis that X_1 and X_2 share a (true) *common* distribution based on the KS distance (D_{max}). However, it is misleading to use this test to indicate if two distributions are *similar*, as it is highly sensitive to large sample sizes, and also as the particular x_1 and x_2 compared here are not strictly independent variables since, for example, nodes with high degrees tend to occur together. Instead D_{max} alone is used in this chapter to indicate the *relative* closeness of distributions.

3.6.2 Kullback-Leibler divergence

The Kullback-Leibler (KL) divergence is also proposed as a suitable metric[7] for comparing network distributions. The KL divergence between two discrete random variables X_1 and X_2 is defined as:

$$D_{KL}(X_1, X_2) = \sum_i P(X_1 = X_i) log \frac{P(X_1=X_i)}{P(X_2=X_i)}$$

where $P(x)$ is the probability of x.

The KL divergence takes account of the difference between the distributions at all points rather than simply at the maximum point. In this chapter, Gaussian kernel density estimation using fixed bins centred around data in the observed data set were found to perform well for as a non-parametric way of estimating the probability density function, although other methods do exist. There are other distance estimation measures also available such as Chi-square statistic, quadratic form distance and match distance which we do not use in this chapter, as most of them rely on the assumption of the underlying sample's distribution.

3.7 Results and Discussion

Most past comparisons of topology generators have been limited to the average node degree, the node degree distribution and the joint degree distribution. The rationale for choosing these metrics is that if those properties are closely reproduced, then the value of other metrics will also be closely reproduced [MKFV06].

In this section I show that current topology generators are able to match first and second order properties well, i.e., average node degree and node degree distribution, but fail to match many other topological metrics. I also discuss the importance of various metrics in the analysis.

3.7.1 Methodology

For each generator I specify the required number of nodes and generate 10 topologies of that size in order to provide confidence intervals for the metrics. I then compute the values of the metrics introduced in Section 3.5 for the generated and observed AS topologies. It is important to note that all topologies studied in this thesis are undirected. Using undirected graphs prevents us from considering peering policies and provider-customer relationships. This is a limitation that is forced upon us by the design of the generators as they do not take such policies into account.

Each topology generator uses several parameters, all of which could be tuned to best fit a particular size topology, e.g., the Skitter dataset. However, there are two problems with attempting this tuning. First, doing so requires selection of an appropriate goodness-of-fit measure of which there are many, e.g., as noted in Section 3.5. Second, in any case tuning parameters to a particular dataset is of questionable merit since, as I argue in Section 3.1, each dataset is only a sample of reality with multiple biases and inaccuracies. Nonetheless, I made a preliminary attempt at tuning in this way for node degree and joint degree distribution in the Waxman model, but it proved of little value with insignificant impact on subsequent results. Consequently, I chose not to pursue this further in this chapter and simply use the

[7] The KL divergence is not strictly a metric as $D_{KL}(X_1, X_2) \neq D_{KL}(X_2, X_1)$

Table 3.1: Comparison of AS level dataset with synthetic topologies.

Topology	Links	Avg. deg.	Max. degree	Top clique size	Max. betweenness	Max. coreness	Assort. coef.	Clust. coef.	Max. closeness
Chinese(n=84)	*211*	*5.02*	*38*	*2*	*1,324*	*5*	*-0.32*	*0.188*	*<0.01*
Waxman	252	6	18	**2**	404	4	0.039	0.117	**0.506**
BA	165	3.93	19	3	**1,096**	2	-0.096	0.073	0.515
GLP	151	3.6	44	3	2,391	**5**	-0.257	**0.119**	0.643
PFP	**250**	**5.95**	**37**	10	849	9	**-0.38**	0.309	0.638
Skitter(n=9204)	*28,959*	*6.3*	*2,070*	*16*	*10,210,533*	*28*	*-0.23*	*0.026*	*<0.01*
Waxman	**27,612**	**6**	33	0	474,673	4	0.205	0.002	**0.264**
BA	18,405	4	190	0	5,918,226	2	-0.05	0.001	0.315
GLP	16,744	3.64	**2,411**	2	34,853,544	5	-0.089	0.003	0.496
INET	18,504	4.02	1,683	3	15,037,631	7	-0.195	0.004	0.514
PFP	27,611	**6**	3,000	**16**	**13,355,194**	**24**	**-0.244**	**0.012**	0.588
RouteViews(n=17446)	*40,805*	*4.7*	*2,498*	*9*	*30,171,051*	*28*	*-0.19*	*0.02*	*<0.01*
Waxman	52,336	6	35	0	1,185,687	4	**0.205**	0.001	**0.25**
BA	34,889	4	392	3	33,178,669	2	-0.04	0.001	0.33
GLP	31,391	3.6	4,226	4	127,547,256	6	-0.08	0.002	0.48
INET	**43,343**	**4.97**	**2,828**	**6**	**31,267,607**	14	-0.258	0.006	0.522
PFP	52,338	6	4,593	23	39,037,735	**30**	-0.252	**0.009**	0.564
UCLA(n=28899)	*116,275*	*8.05*	*4,393*	*10*	*76,882,795*	*73*	*-0.165*	*0.05*	*0.32*
Waxman	86,697	6	40	0	3,384,114	4	0.213	<0.001	0.246
BA	57,795	4	347	0	52,023,288	2	-003	<0.001	**0.3**
GLP	52,456	3.63	7391	2	371,651,147	6	-0.08	<0.001	0.486
INET	**91,052**	**6.3**	**6,537**	**12**	**88,052,316**	38	-0.3	**0.01**	0.55
PFP	86,696	6	8076	26	123,490,676	**40**	**-0.218**	**0.01**	0.57

default values embedded within each generator. This corresponds to the way in which such generators are generally used. I address the problem of parameter optimisation in Chapter 4.

3.7.2 Topological metrics

In this section I discuss the results for each metric separately and analyse the reasons for differences between the observed and the generated topologies.

Table 3.1 displays the values of various metrics (columns) computed for different topologies (rows). Blocks of rows correspond to a single observed topology and the generated topologies with the same number of nodes as the observed topology. Bold numbers represent nearest match of a metric value to that for the relevant observed topology. Rows in each block are ordered with the observed topology first followed by the generated topologies from oldest to newest generator. Each metric's value is the calculated value for the observed topology, and the average of the 10 synthetic topologies for each generator. Note that Inet requires the number of nodes to be greater than 3037 and hence cannot be compared to the Chinese topology.

I observe a small but measurable improvement from older to newer generators in how well they match some measures such as maximum degree, maximum coreness, and assortativity coefficient. This suggests that topology generators have been successively improved to better match some properties of the observed topologies. However, the number of links in the generated topologies may differ considerably from the observed topology due to the assumptions made by the generators.

Waxman and BA models fail to capture the maximum degree, the top clique size, maximum be-

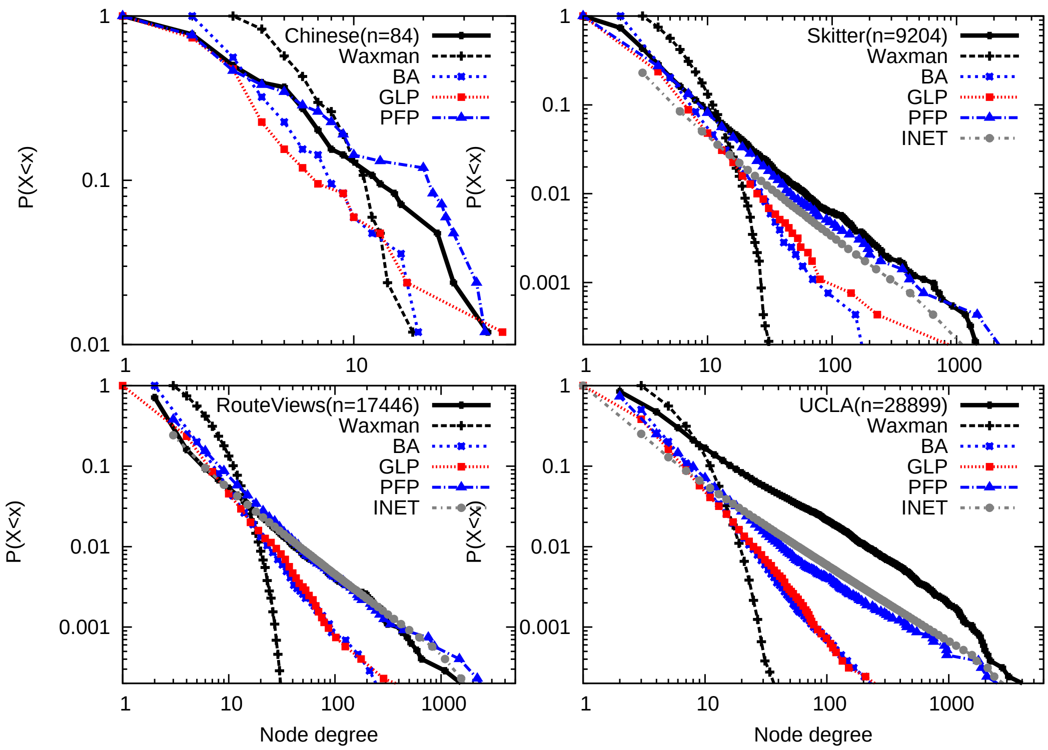

Figure 3.2: Comparison of node degree CCDFs.

tweenness and coreness. Those two generators are too simplistic in the assumptions they make about the connectivity of the graphs to generate realistic AS topologies. Waxman relies on a random graph model which cannot capture the clique that is known to exist between tier-1 ASes, nor the heavy tail of the node degree distribution. BA tries to reproduce the power law node degrees with its preferential attachment model but fails to reach the maximum node degree by far as it only adds edges between new nodes and not between existing ones. Hence neither of these two models is able to create the highly-connected core of tier-1 ASes.

PFP and Inet manage to come closer to the values of the metrics of the observed topologies. For Inet this is due to the fact that nodes are fully meshed (at the core), whereas for PFP it is its rich-club connectivity model that allows it to add edges between existing nodes. Based on the observations, I conclude that the core of the Internet is tending towards a fully meshed network.

Node degree distribution

Figure 3.2 shows the CCDF of the node degree for all topologies on a log-log scale. We observe that the Chinese topology does not exhibit power law scaling due to its limited size, whereas all the larger AS topologies do exhibit power law scaling of node degrees. The Waxman generator completely fails to capture this behaviour as it is based on a random-graph model, but recent topology generators do capture this power law behaviour of the node degrees quite well. In the case of the RouteViews and UCLA datasets, Inet and PFP outperform other topology generators. Note that, contrary to RouteViews where the degree distribution displays strict power law scaling, the UCLA dataset has a slightly concave shape. In summary, more recent generation models reproduce node degree distribution well, as expected since

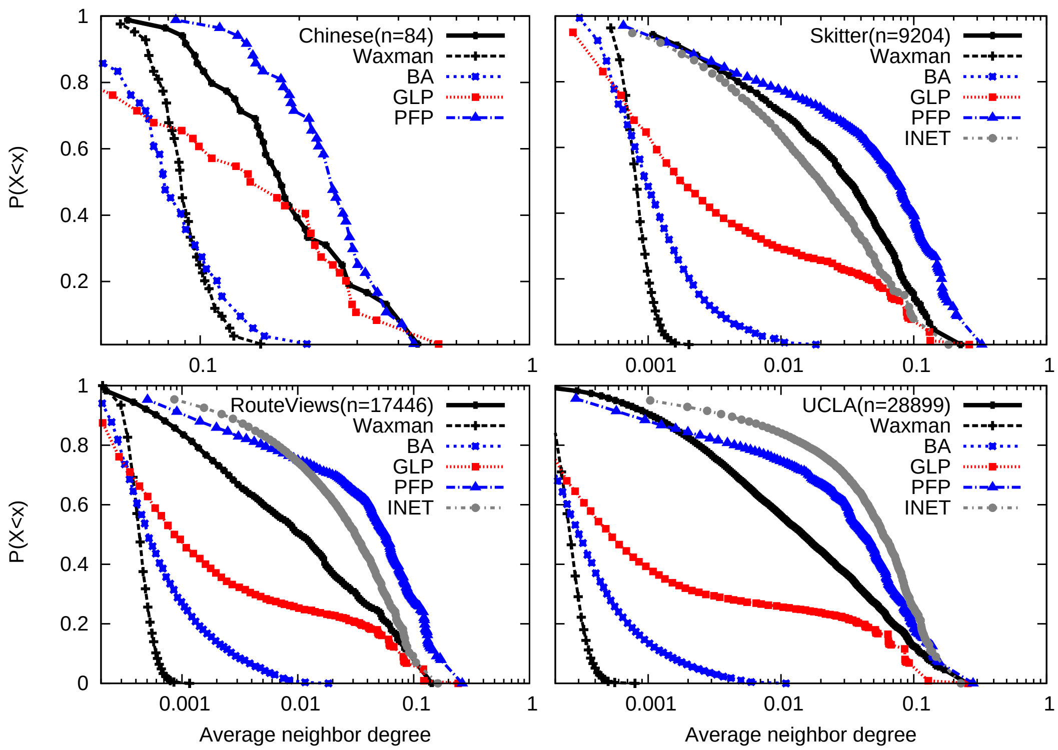

Figure 3.3: Comparison of average neighbour connectivity CCDFs.

most focus has been on this metric.

Average neighbour connectivity

Neighbour connectivity has been far less studied than node degree, although it is very important to match local interconnection among a node's neighbours when reproducing the topological structure of the Internet [MKF+06]. Figure 3.3 shows the CCDF of the average neighbour degrees for all topologies. We observe that Waxman, BA and GLP all underestimate the local interconnection structures around nodes due to their simplistic way of modelling node interconnections. Note that BA and GLP typically generate graphs with far fewer links than the observed topologies so they underestimate neighbour degrees on average. For the larger topologies, i.e., RouteViews and UCLA, PFP and Inet typically overestimate the neighbour connectivity, as they both place a large number of inter-As links at the core. In addition, the shapes of the neighbour connectivity CCDF differ for the larger topologies: Inet and PFP have two regimes, one for high-degree nodes, and another for low-degree nodes. On the other hand, observed topologies have a smooth region for the high-degree nodes, followed by a rather stable region which caused by similar degree nodes. We observe that the highest degree nodes in the UCLA topology have very high values of neighbour connectivity. This is consistent with the belief that tier-1 providers are densely meshed. In summary, existing topology generators do not reproduce local interconnection behaviour well, but it is an important aspect of today's AS topology and may significantly alter the quality of results from simulations relying on the AS topology.

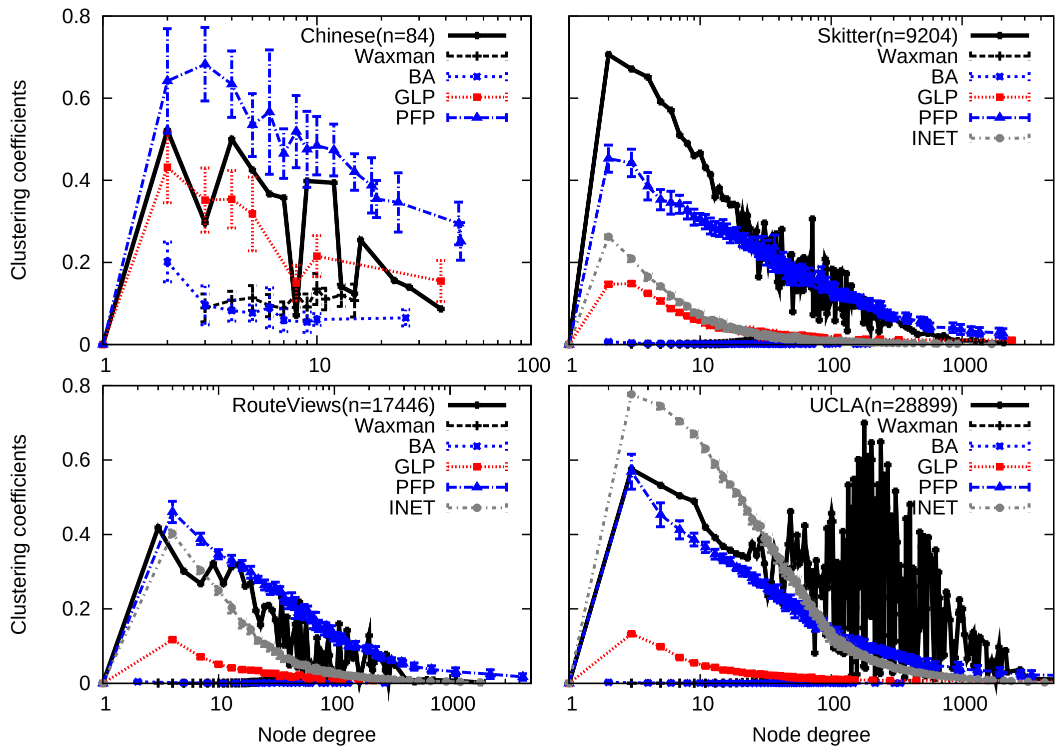

Figure 3.4: Comparison of clustering coefficients.

Clustering coefficients

Like the average neighbour connectivity, the clustering coefficient gives information about local connectivity of the nodes. It is important to reproduce clustering due to its impact on the local robustness in the graph: nodes with higher local clustering have increasing local path diversity [MKF+06]. Clustering properties of a graph can directly affect simulation on performance of multipath and resilience of overlay routing.

Figure 3.4 displays the clustering coefficients of all nodes in the topologies. Error bars indicate 95% confidence intervals around the mean values of the 10 topologies from each generator. We observe that Waxman and BA significantly underestimate clustering, which is again consistent with their simplistic way of connecting nodes. GLP approximates the clustering of the Chinese topology quite well but fails in the case of the larger observed topologies. PFP and Inet capture clustering reasonably well compared to the other topology generators. However, Inet does not reproduce the tail of the distribution well due to the randomness factor in its model for edge addition once the core is fully meshed.

We also observe that for medium degree nodes, clustering coefficients display rather high variability which increases with the size of the observed topologies. This behaviour seems to be a property of the observed AS topology of the Internet (Section 3.8), and not just an artifact of the incompleteness of observed AS topologies.

In summary, all topology generators fail to properly capture the clustering of the AS topology. Generators typically underestimate the local connectivity. Only Inet for the UCLA topology overestimates connectivity of low-degree nodes while underestimates it for high-degree nodes. The current topology

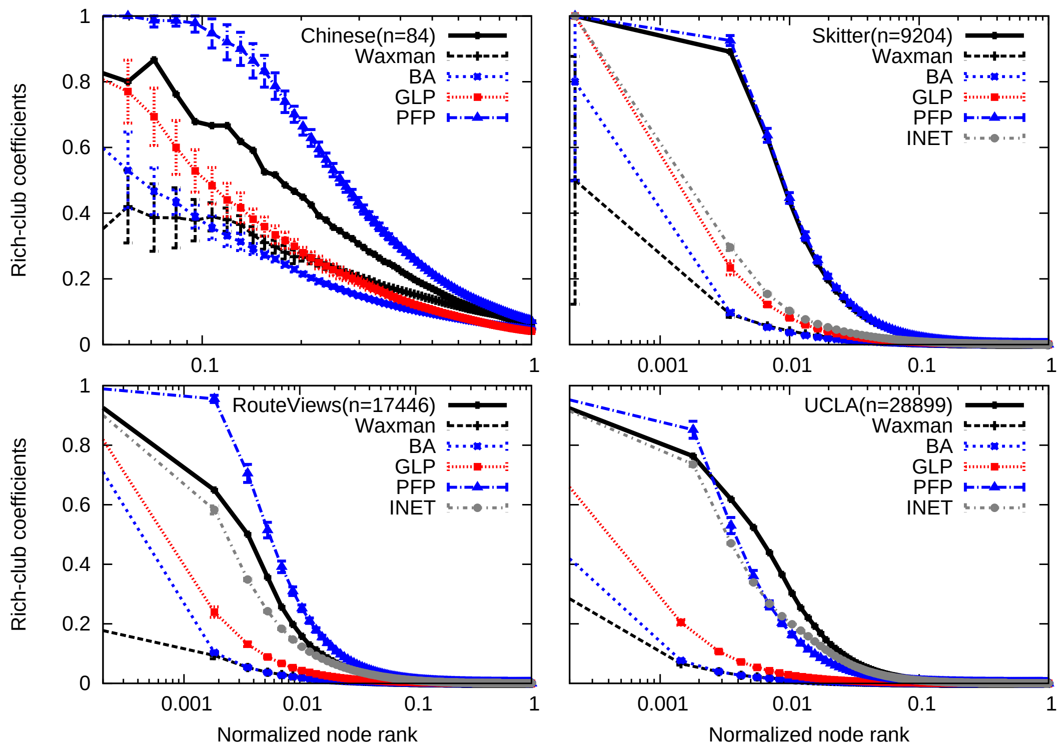

Figure 3.5: Comparison of rich-club connectivity coefficients

generators do not seem to have a proper model of local node connectivity.

Rich-club connectivity

Rich-club connectivity gives information about how well-connected among themselves are the nodes of high degree. Figure 3.5 makes it clear that the cores of the observed topologies are very close to a full mesh, with values close to 1 on the left of the graphs. The error bars again indicate the 95% confidence intervals around the mean values of the different instances of the generated topologies. Waxman and BA perform poorly for this measure in general. Only PFP and Inet generate topologies with a dense enough core compared to the observed topologies. However, PFP overestimates the rich-club connectivity of the Chinese and RouteViews topologies which is consistent with the emphasis that PFP gives to the rich-club connectivity in its design. Inet performs well due to its emphasis on a highly connected core, especially for larger topologies where data has been collected across multiple peering points.

In summary, most topology generators underestimate the importance of rich-club connectivity of the AS topology. PFP is the only topology generator that emphasises the importance of the dense core of the AS topology.

Shortest path distributions

Figure 3.6 displays the distributions of shortest path length. Apart from BA, most topology generators approximate the shortest path length distribution of the Chinese graph quite well due to its small size and thus limited scope for error. For the other topologies, PFP and Inet generally underestimate the path length distribution while Waxman and BA overestimate. Particular generators seem to capture the path length distribution for particular topologies well: PFP matches that for Skitter well and GLP is

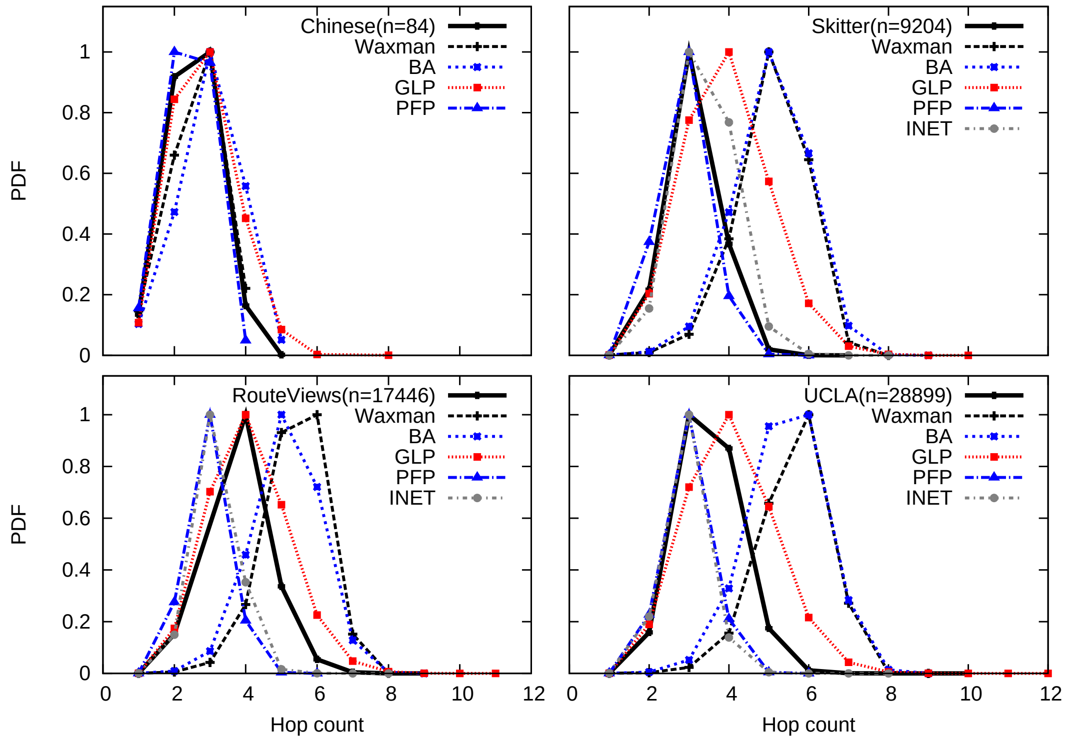

Figure 3.6: Comparison of shortest path distributions (number of hops).

close for Routeviews. Inet and PFP both do a better job for UCLA than for RouteViews but both still underestimate the distribution.

In summary, shortest path length is not well captured by any topology generator. Given the poor match of generators on local connectivity metrics, it is not surprising.

Spectrum

The spectrum of the normalised Laplacian matrix is a powerful tool for characterising properties of a graph. If two large graphs have the same spectrum, they have the same topological structure.

Figure 3.7 displays the CDF of the eigenvalues computed from the normalised Laplacian matrix of each topology.

As with other topological metrics, Inet and PFP perform best. The difference between the topology generators is most easily observed around the eigenvalues equal to 1. These eigenvalues play a special role as they indicate repeated duplications of topological patterns within the network. By duplication, I mean different nodes having the same set of neighbours giving their induced subgraphs the same structure. Through repeated duplication, one can create networks with eigenvalue 1 of very high multiplicity [BJ07]. In addition, we observe that the spectra have a high degree of symmetry around the eigenvalue 1. If a network is bipartite, i.e., it consists of two connected parts with no links between nodes of the same part, then its spectrum will be symmetric about 1. Consequently, the observed AS topologies appear close in spectral terms to a bipartite graph, another phenomenon that arises through repeated structure duplication. In the AS topology many ASes share a similar set of upstream ASes without being directly connected to each other. Inet and PFP are good examples of topology generators

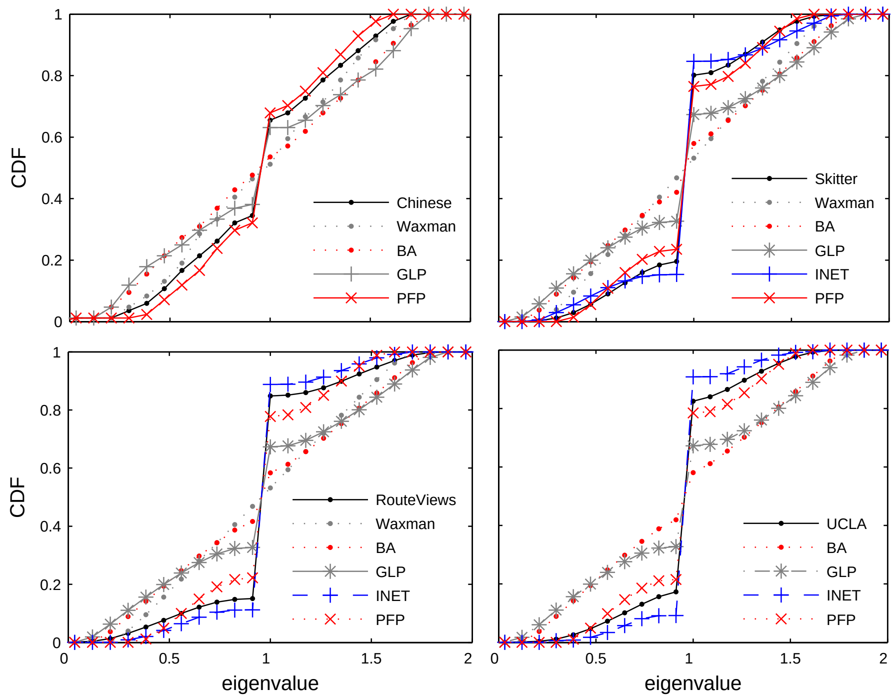

Figure 3.7: Comparison of cumulative distributions of eigenvalues (from normalised laplacian).

where this strategy is implemented. Note that the simple preferential attachment model of BA does not reproduce the eigenvalues around 1 very well. In the simple BA model, new nodes connect randomly to a given number of existing nodes, favouring connections to high degree nodes. In the Internet in contrast, although small ASes may tend to connect to large upstream providers, they might not connect preferentially to the largest ones, connecting instead to national or regional providers. In summary, these results provide further evidence that the interconnection structure of the AS topology is more complex than current models assume.

3.7.3 Measures of similarity

In Section 3.7.2, I presented visual evidence for the (dis)similarity both among topology generators and between generators and observed topologies. In this section I present a more objective approach, based on the statistical distance measures described in Section 3.6: the Kolmogorov-Smirnov (KS) distance and the Kullback-Leibler (KL) divergence.

In the following tables, the values of the distances and the standard deviations are shown for the topological metrics with distributions: node degree, neighbour connectivity, clustering coefficient, and rich-club coefficient. I provide the average values of the statistical distances and the standard deviation around the average over the 10 topologies generated by each topology generator. When no deviation is shown, it was < 0.01.

Both statistical measures globally confirm the visual inspection of Section 3.7.2: more recent topology generators produce topologies whose properties are closer to the observed topologies. Table 3.2

Table 3.2: Statistical distances for Chinese vs. synthetic topologies.

	Node degree		Neighbour connectivity	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.27±0.07	**0.6±0.1**	0.75±0.03	27.4±4.1
BA	**0.12±0.03**	3.5±1.8	0.74±0.07	18.4±8.1
GLP	0.24±0.08	0.64±0.31	**0.41±0.08**	1.18±0.72
PFP	0.17±0.04	1.45±0.48	0.51±0.07	**0.85±0.25**
	Clus. Coefficients		Rich-Club Coefficients	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.61±0.03	22.31±4.5	0.22±3.5	4.2±2.8
BA	0.65±0.1	13.5±5.2	0.28±0.01	2.78±1.4
GLP	**0.31±0.05**	1.08±.6	0.26±0.04	0.34±0.19
PFP	0.32±0.11	**0.34±0.14**	**0.12±0.01**	**0.11±0.02**

Table 3.3: Statistical distances for Skitter vs. synthetic topologies.

	Node degree		Neighbour connectivity	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.54±0.04	2.27±0.15	0.99±0.01	44.48±0.08
BA	0.41±0.02	17.1±2.6	0.99±0.01	44.7±0.25
GLP	0.31±0.06	17.42±4.1	0.31	2.16
Inet	**0.075±0.02**	**4.13**	0.40±0.02	**1.82±0.31**
PFP	0.13±0.03	18.2±2.31	**0.13±0.05**	18.2±2.21
	Clust. Coefficients		Rich-Club Coefficients	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.91±0.02	40.62±1.2	0.2±0.05	6.75±1.3
BA	0.9±0.05	44.62±0.12	0.37±0.09	7.34±1.21
GLP	**0.7±0.02**	19.12±1.8	0.3±0.01	4.34±.45
INET	0.74±0.01	11.34±1.23	0.25	3.82±0.2
PFP	0.09±0.02	**0.59±0.19**	**0.03**	**0.91±0.14**

provides the KS and KL results for topology generators against the Chinese topology for the four chosen topological metrics. Topology generators do not show improvement for the node degree. However, for the other three metrics successive topology generators do show improvement. Overall, the PFP and GLP model both have small relative distances to the Chinese dataset, due to the small size of the dataset, the presence of high degree nodes as core ASes and fewer inter-AS connections.

Table 3.3 displays the results of the statistical measures for results against the Skitter topology. We observe a particularly good match of the node degree distribution by Inet. PFP outperforms all other topology generators for the clustering coefficients and the rich-club coefficients, consistent with the visual inspection.

Statistical distances for RouteViews (Table 3.4) show that Inet again better matches the node degree distribution. GLP and Inet both perform better than other generators for neighbour connectivity. PFP performs better than the others on the clustering coefficients. On the other hand, none of the generators manages to obtain a close distance for the rich-club coefficients. On Figure 3.5, Inet seemed to be close to

Table 3.4: Statistical distances for RouteViews vs. synthetic topologies.

	Node degree		Neighbour connectivity	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.5±0.03	50.77±0.01	0.94±0.01	42.68±0.25
BA	0.2±0.02	50.74±0.01	0.94±0.01	42.91±0.34
GLP	0.18±0.03	50.73±0.01	**0.12±0.02**	**0.1±0.02**
Inet	**0.07**	**9.92**	0.23±0.02	0.2±0.01
PFP	0.11±0.03	50.7	0.62±0.02	1.25±0.07
	Clust. Coefficients		Rich-Club Coefficients	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.83±0.05	39.4±1.2	0.97	42.23±0.43
BA	0.96±0.01	44.08±0.21	0.97	43.07±0.6
GLP	0.58±0.02	12.9±0.65	0.96	40.7±0.9
INET	0.39±0.01	1.35±0.2	0.93	34.18±1.1
PFP	**0.32±0.06**	**0.21±0.03**	**0.92**	**27.4±2.45**

Table 3.5: Statistical distances for UCLA vs. synthetic topologies.

	Node degree		Neighbour connectivity	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.52±0.01	**1.33±0.9**	0.99±0.01	46.31±1.3
BA	0.17±0.03	2.15±0.8	0.99±0.01	46.42±0.7
GLP	0.18±0.05	2.21±0.7	0.32±0.03	0.63±0.04
Inet	0.2±0.02	5.34	**0.29±0.01**	**0.41±0.01**
PFP	**0.12±0.03**	2.17±0.8	0.48±0.05	0.83±0.21
	Clust. Coefficients		Rich-Club Coefficients	
	KS distance	KL divergence	KS distance	KL divergence
Waxman	0.93±0.02	44.2±0.34	0.31	14.5±4.32
BA	0.99±0.01	45.42	0.5	14.32±2.3
GLP	0.82±0.01	33.32±0.9	0.42±0.01	8.9±1.2
INET	**0.38±0.01**	**0.53±0.01**	**0.13**	**2.85±0.12**
PFP	0.38±0.02	0.79±0.15	0.16	3.23±0.4

RouteViews for rich-club coefficients, but this is not supported by the statistical distances. The behaviour for rich-club connectivity is surprising, especially for PFP which is highly biased towards reproducing rich-club connectivity. I believe this is due mainly to the addition of many extra peering links in this dataset, which was not captured by model designers.

Statistical tests results for UCLA (Table 3.5) reveal a more complex picture. For node degrees, no generator seems to outperform the others, although Inet performs worst. GLP, Inet and PFP perform equally well on the neighbour connectivity. For clustering coefficients and rich-club connectivity, Inet and PFP perform better than the others.

Visual inspection of Section 3.7.2 seemed to suggest that each successive topology generator introduced improvements in their matching of observed AS topologies. Waxman and BA perform poorly both in visual inspection and in the statistical distances. The KL divergences clarify the difference of the two distributions across all the values and hence minimise the effects of local differences at certain

Table 3.6: Comparison of AS topology datasets from multiple peering points.

Topology	Nodes	Links	Avg. deg.	Max. degree	Top clique size	Max. betweenness	Max. coreness	Assort. coef.	Clust. coef.	Max. closeness
1 peer	17,952	34,617	3.86	980	4	35,069,182	9	-0.18	0.008	$<$0.01
10 peers	27,838	64,717	4.65	2,731	7	52,862,315	20	-0.18	0.007	$<$0.01
25 peers	27,885	67,659	4.85	2,808	7	49,798,002	25	-0.19	0.01	$<$0.01
All peers	27,924	70,064	5.02	3,371	7	70,142,726	30	-0.18	0.01	$<$0.01

values.

The statistical measures show that apparent visual closeness of two distributions does not mean close distance in distributional terms, due partly to the use of logarithmic scale axes. Improvements in successive topology generators are not consistent across all metrics and across all observed topologies. Nonetheless, most of the time the most recent generators, Inet and PFP, do outperform the other topology generators. This indicates that more attention should be given on capturing the effects of peering links in the core and at the edge of the AS topology, as this is the significant difference between these two generators and the older Waxman and BA generators.

3.8 Multiple Vantage Points

The previous section studied in detail how well topology generators capture the properties of observed AS topologies. In this section, I will argue about why topology generators capture different properties of observed AS topologies with varying degrees of success. To that end I examine the impact on the metrics of the number of vantage points from which BGP data is collected. For the analysis I used collected BGP data from over 40 RouteViews peering points, for a period of 6 months from May 2007. This time period was chosen to be the same as that used to build the UCLA dataset.

Table 3.6 shows the values of the topological metrics the same way as in Table 3.1, for AS topologies obtained from different numbers of observation points. When comparing the AS topologies using 1 (average value amongst all peers) and 10 random observation points, we see a significant increase in the number of nodes and links. Hence, one might also expect a significant difference in the other metrics, and indeed, the maximum node degree almost triples and the number of fully-meshed nodes almost doubles. As a consequence, the size of the core increases as indicated by the maximum coreness value. In turn, the number of shortest paths crossing the core also increases as indicated by the maximum betweenness. On the other hand, we see that going from 1 to 10 observation points slightly decreases the value of the clustering coefficient. Most probably this is because with 10 observation points we discover more of the core than the edge of the network, which does not contribute to increase the overall value of the clustering coefficient. With 25 or more observation points the links on the edge of the network are also discovered more, contributing to the increase of the value of the clustering coefficient. This behaviour is confirmed by a slight decrease of the value of the maximum betweenness from 10 to 25 observation points.

Preferential attachment models originate in the belief that small ASes tend to connect to large upstream ASes, leading to a disassortative network. Although the value of the assortativity coefficient

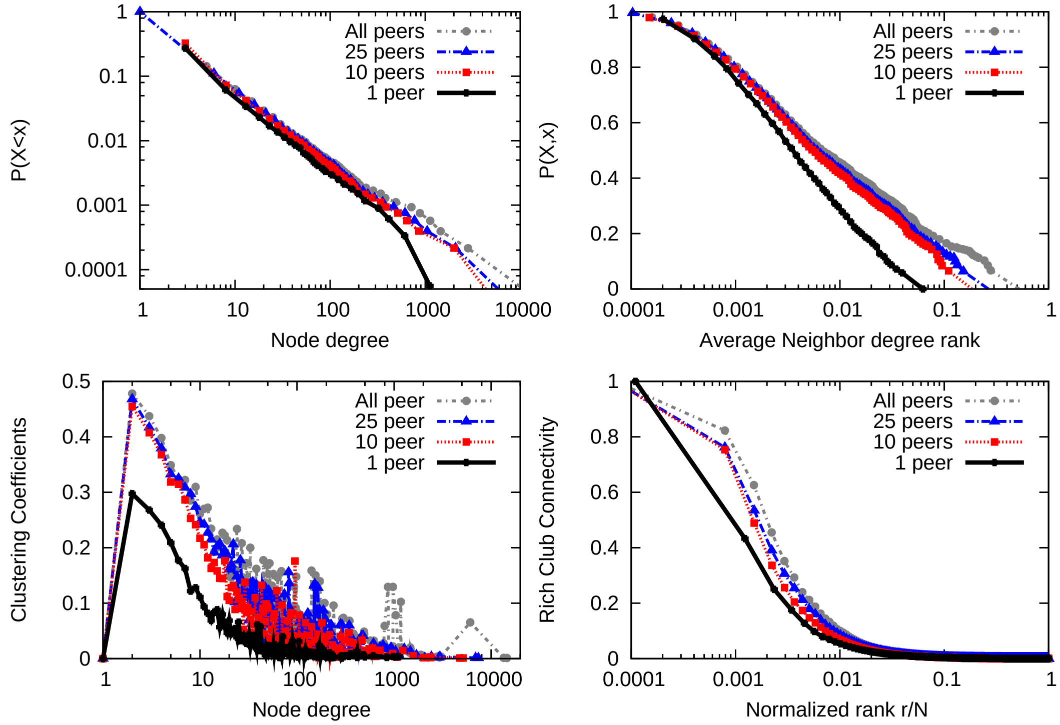

Figure 3.8: Comparison of effects of the number of peering points.

Recent work [RTM08] estimates that more than 700 observations may be needed in order to discover nearly all missing links. However even this figure is an estimate and may not able to find the ground truth.

is negative for the AS topology, it is not affected by an increase in the number of observation points. The links added by increasing the number of observation points seem to be neutral for the assortativity of the AS topology. One implication is that the links that can be discovered by using more observation points do not preferentially interconnect ASes of any particular degree. I conjecture that this is due to the type of peering relationships that are missed. If node degrees give an indication of the likely type of peering relationship, then I suggest that BGP does not preferentially miss peer-peer relationships, which are believed to be more difficult to observe that customer-provider ones due to the nature of BGP path advertisements [CGJ+04].

I now turn in more detail to the effect of the number of peering points on four particular topological metrics (see Figure 3.8). The addition of observation points mostly affects node degree distribution for high degree nodes. As I increase the number of observation points, on average the neighbours of a node will have a higher degree. However, this does not hold for nodes whose neighbours already have high degrees (left part of the average neighbour degree curves). Those nodes correspond to stub networks connected to very well interconnected upstream providers. For the clustering coefficient, when moving from one to several observation points, the difference is striking. For all node degrees, the clustering coefficient significantly increases. On the other hand, when moving from a few peerings to many, the difference appears most for high degree nodes. This illustrates the better observability of links in the core compared to the edge of the network. Rich-club connectivity confirms the previous observations in that adding a few observation points is enough to discover the core links.

In this section I have illustrated the importance of relying on a sufficiently large number of observation points in order to properly capture the actual properties of the AS topology. Using only a few observation points has led researchers to simplify the complexity of the interconnection structure between ASes. The improper AS topology on which researchers have relied has caused the creation of topology generators that underestimate this interconnection structure between ASes. The results show that researchers must use rich datasets for an accurate understanding of the Internet AS topology.

3.9 Conclusions and Contributions

In this chapter, I provide insight into the Internet's AS topology. I compare multiple synthetic topologies from generators based on different models, both among themselves and to several observed AS topologies collected at different times using different methods. I base this comparison on numerous topological metrics, and use statistical measures to perform this comparison objectively.

My analysis revealed that current topology models do not faithfully represent the reality of the Internet AS topology. Current models over-emphasise node degree distribution and preferential attachment, while failing to reproduce local connectivity metrics. Although I observe that more recent topology generators generally perform better than older ones, I find that metrics giving information about local connectivity properties were not well captured by any existing topology generator. In addition to clustering and centrality properties, the highly meshed core of the Internet AS topology must be considered in order to generate representative synthetic topologies, increasing the quality of simulations based upon them.

I also compared the properties of AS topologies relying on different sets of observations. I observed that, in contrast to structural metrics, node degree-related properties are not greatly affected by the addition of more vantage points as they add only a small percentage of peering links. On the other hand, the power-law nature of the node degree distribution seems questionable, as increasing the number of observation points causes deviation from strict power-law scaling.

Finally, I wish to point out that the AS topology, useful as it is, provides only limited information about the Internet's size and other properties. When creating AS topologies, not all ASes should be considered equal. Some networks may contain thousands of routers and links and be represented by a single AS number, whereas others may have their own AS number but contain just a single router. Future AS topology generators should permit the addition of metadata such as peering relationship and relative importance of nodes[8].

[8] The work presented in this chapter is the result of collaboration with Damien Fay, Steve Uhlig, Olaf Maennel and my advisors. Damien Fay was mainly responsible for the accuracy of use of statistical measures. Steve Uhlig contributed to the use of spectrum. Olaf Maennel provided the BGP data. All authors collaborated on the writing. However, the largest part of the underlying ideas behind the work and methodological comparison approach, all the code and the detailed analysis of the collected traces have been done by me.

Chapter 4

Tuning Topology Generators

An increasing number of synthetic topology generators are available, each claiming to produce representative Internet topologies. Every generator has its own parameters, allowing the user to generate topologies with different characteristics. However, there exist no clear guidelines on tuning the value of these parameters in order to obtain a topology with specific characteristics.

In this chapter I tune the parameters of several topology generators to match a given Internet topology. The optimisation is performed either with respect to the link density, or to the spectrum of the normalised Laplacian matrix. Contrary to approaches in the literature that rely only on the largest eigenvalues, I take into account the distribution of eigenvalues. However, I show that on their own the eigenvalues cannot be used to construct a metric for optimising parameters. Instead I use a weighted spectral method which simultaneously takes into account all the properties of the graph.

4.1 Introduction

Today's Internet is formed from more than 25,000 ASes, each of which can contain few or hundreds of routers. Constant evolution and change in the Internet, due to failures and router configuration bugs in the short term, and growth and death of networks in the long term, has made it difficult for scientists to produce representative Internet topologies at either AS or router level. However, such maps are essential for the simulation and analysis of ideas including new and improved routing protocols, and peer-to-peer or media-streaming applications. Since obtaining accurate, timely maps of the Internet topology is difficult, and development of new protocols and systems requires understanding their performance over a range of scenarios, researchers use synthetic topology generators.

There are many such generators, each of which is parameterised, often with multiple parameters, giving rise to a plethora of potential synthetic graphs. Understanding and generating those graphs, useful because they accurately represent features of the true underlying Internet graph, is difficult. Existing approaches to tuning the generator parameters range from selection of particular metrics of interest, e.g., link count, and tuning to match that particular metric, to simply using the default parameters encoded in the particular release of the generator package in use.

The core problem is to select an appropriate cost function which reflects those aspects of the graph that are important to the user and weights those aspects accordingly. Such a selection process is inherently subjective: there is no “best” cost function in general. Once a suitable cost function is selected, it is a simple matter to tune the available parameters of the topology generator to produce output that optimally matches said cost function.

In the light of this, the contributions in this chapter are as follows:

- I propose a new cost function, the *weighted spectrum*, constructed from the eigenvalues of the normalised Laplacian matrix, or graph spectrum;
- I demonstrate that the graph spectrum alone is unsatisfactory as a cost function;
- I use an efficient approximation of the weighted spectrum which favours the more significant eigenvalues;
- I use this approximation to tune parameters for a set of Internet topology generators, enabling us to use these generators to effectively match a particular measured Internet topology.

The graph spectrum is a useful starting point for such a cost function as it yields a set of invariants about a graph that encode all the properties of that graph [Chu97]. The proposed cost function improves on the simple graph spectrum because it incorporates the knowledge that not all eigenvalues are equally important, and weights toward those that are considered to encode more significant aspects of the graph’s structure. The basis of the algorithm is to provide a way to measure the difference between two graphs with respect to a common reference, a suitable regular graph.[1]

After reviewing related work in Section 4.2, I outline background theory in Section 4.3. In Section 4.4 I present the results of the analysis and in Section 4.5 I compare topologies generated at optimal values of the parameters with an observed dataset. Finally, I conclude the chapter in Section 4.6 and discuss future work.

4.2 Related Work

Zegura *et al.* [ZCD97] analyse topologies of 100 nodes generated using pure random, Waxman [Wax88], exponential and several locality based models of topology such as Transit-Stub [CDZ97]. They use metrics such as average node degree, network diameter, and number of paths between nodes, and use the number of edges as the metric of choice for optimisation of the tuning parameter. However as I show in this chapter, the number of links is not an ideal choice particularly in random networks, due to the network structure only resembling the observed Internet topology at link counts much higher than those suggested by the optimisation process.

Tangmunarunkit *et al.* [TGJ^{+}02] provide a first point of comparison of the underlying characteristics of degree-based models against structural models. A major conclusion is that the degree-based model in its simplest form performs better than random or structural models at representing all the studied parameters. They compare three categories of model generators: the Waxman model of random graphs, the TIERS [Doa96] and Transit-Stub structural models, and the simplest degree based generator, called the Power-Law Random Graph [ACL00]. They compare under three metrics: expansion,

[1] A regular graph is one where all nodes have the same degree.

resilience and distortion and conclude that the hierarchy present in the measured networks is more strict than in degree-based generators. However, they leave many questions unanswered about the accuracy of degree-based generators and their choice of metrics and parameter values.

Heckmann *et al.* [HPSS03] discuss different types of topologies and present a collection of real-world topologies that can be used for simulation. They then define several similarity metrics, such as the shortest path distributions, node degree distributions and node rank exponents, to compare artificially generated topologies with real world topologies from AT&T's network. They use these to determine the input parameter range of the topology generators of BRITE [MLMB01], TIERS and GT-ITM [CDZ97] to create realistic topologies.

Gkantsidis *et al.* [GMZ03] perform a comparison of clustering coefficients using the eigenvectors of the k largest eigenvalues of the adjacency matrices of BGP topology graphs. However, the selected eigenvectors are all given equal importance. They do not take into account the rest of the spectrum, although it has recently been shown that the eigenvalues of either the adjacency matrix or the normalised Laplacian matrix can be used to accurately represent a topology and some specific eigenvalues provide a measure of properties such as robustness of a network to failures [But06, JU07].

Vukadinovic *et al.* [VHE02] used the normalised Laplacian spectrum for analysis of AS graphs. They propose that the normalised Laplacian spectrum can be used as a fingerprint for Internet-like graphs. Using the Inet [WJ02] generator and AS graphs from BGP data, they obtain eigenvalues of the normalised Laplacian matrix. The differences between synthetic and observed topologies indicate that the structural properties of the Internet should be included in an Internet AS model alongside power law relationships. They believe that the graph spectrum should be considered an essential metric when comparing graphs. I expand on this work by demonstrating how an appropriate weighting of the eigenvalues can be used to reveal structural differences between two topologies.

Use of spectrum for graph comparison is not limited to Internet research. Hanna [Han07] uses graph spectra for numerical comparison of architectural space in large building plans. By defining space as a graph, he shows that the spectra of two plan types can be used effectively to judge the effects of global vs. local changes to, and hence the edit distances between, the plans. Hanna believes spectra give a reliable metric for capturing the local relationships and can be used to guide optimisation algorithms for reproducing plans.

4.3 Weighted Spectral Distribution

I use the Weighted Spectral Distribution (WSD), which is related to another common structural metric, the clustering coefficient, for examining the characteristics of networks with different mixing properties.

Denote an undirected graph as $G = (V, E)$ where V is the set of vertices (nodes) and E is the set of edges (links). The adjacency matrix of G, $A(G)$, has an entry of one if two nodes, u and v, are connected and zero otherwise

$$A(G)(u,v) = \begin{cases} 1, & \text{if } u,v \text{ are connected} \\ 0, & \text{if } u,v \text{ are not connected} \end{cases} \tag{4.1}$$

Let d_v be the degree of node v and $D = diag(sum(A))$ be the diagonal matrix having the degrees along its diagonal. Denoting by I the identity matrix $(I)_{i,j} = 1$ if $i = j, 0$ otherwise, the Normalised Laplacian L associated with graph G is constructed from A by normalising the entries of A by the node degrees of A as

$$L(G) = I - D^{-1/2}AD^{-1/2} \tag{4.2}$$

or equivalently

$$L(G)(u, v) = \begin{cases} 1, & \text{if } u = v \text{ and } d_v \neq 0 \\ -\frac{1}{\sqrt{d_u d_v}}, & \text{if } u \text{ and } v \text{ are adjacent} \\ 0, & \text{otherwise} \end{cases} \tag{4.3}$$

As L is a real symmetric matrix there is an orthonormal basis of real eigenvectors $e_0, \ldots, e_{n-1}$ (i.e., $e_i e_j^T = 0$ and $e_i e_i^T = 1$) with associated eigenvalues $\lambda_0, \ldots, \lambda_{n-1}$. It is convenient to label these so that $\lambda_0 \leq \ldots \leq \lambda_{n-1}$. The set of pairs (eigenvectors and eigenvalues of L) is called the spectrum of the graph. It can be seen that

$$L(G) = \sum_i \lambda_i e_i e_i^T \tag{4.4}$$

The eigenvalues $\lambda_0, \ldots, \lambda_{n-1}$ represent the strength of projection of the matrix onto the basis elements. This may be viewed from a statistical point of view [SR03] where each $\lambda_i e_i e_i^T$ may be used to approximate $A(G)$ with approximation error inversely proportional to $1 - \lambda_i$. However, for a graph, those nodes which are best approximated by $\lambda_i e_i e_i^T$ in fact form a cluster of nodes. This is the basis for spectral clustering, a technique which uses the eigenvectors of L to perform clustering of a dataset or graph [NLCK05]. The first (smallest) non-zero eigenvalue and associated eigenvector are associated with the main clusters of data. Subsequent eigenvalues and eigenvectors can be associated with cluster splitting and also identification of smaller clusters [NJW02]. Typically, there exists what is called a *spectral gap* in which for some k and j, $\lambda_k \ll \lambda_{k+1} \approx 1 \approx \lambda_{j-1} \ll \lambda_j$. That is, eigenvalues $\lambda_{k+1}, \ldots, \lambda_{j-1}$[2] are approximately equal to one and are likely to represent noise in the original dataset, i.e., links in a graph which do not belong to any particular cluster. It is then usual to reduce the dimensionality of the data using an approximation based on the spectral decomposition. However, in this work I am interested in representing the global structure of a graph (e.g. I am interested in the presence of many small clusters), which is essentially the spread of clustering across the graph. This information is contained in all the eigenvalues of the spectral decomposition.

A full derivation of WSD is present in [HFU+08]. To summarise: the eigenvalues of L lie in the range 0 to 2 (the smallest being 0), i.e., $0 = \lambda_0 \leq \ldots \leq \lambda_{n-1} \leq 2$, and their mean is 1.

The distribution of the n numbers $\lambda_0, \ldots, \lambda_{n-1}$ contains useful information about the network, as will be seen. In turn, information about this distribution is given by its moments in the statistical sense, where the N^{th} moment is $1/n \sum_i (1 - \lambda_i)^N$. These moments have a direct physical interpretation in terms of the network, as follows. Writing B for the matrix $D^{-1/2}AD^{-1/2}$, so that $L = I - B$, then

[2] i.e., the eigenvalues at the centre of the spectrum.

by (4.3) the entries of B are given by

$$(D^{-1/2}AD^{-1/2})_{i,j} = \frac{A_{i,j}}{\sqrt{d_i}\sqrt{d_j}} \tag{4.5}$$

Now the numbers $1-\lambda_i$ are the eigenvalues of $B = I - L$, and so $\sum_i (1-\lambda_i)^N$ is just $\mathrm{tr}(B^N)$.[3] Writing $b_{i,j}$ for the (i,j)-th entry of B, the (i,j)-th entry of B^N is the sum of all products $b_{i_0,i_1} b_{i_1,i_2} \ldots b_{i_{N-1} i_N}$ where $i_0 = i$ and $i_N = j$. But $b_{i,j}$, as given by (4.5), is zero unless nodes i and j are adjacent. So we define an N-cycle in G to be a sequence of vertices $u_1 u_2 \ldots u_N$ with u_i adjacent to u_{i+1} for $i = 1, \ldots, N-1$ and with u_N adjacent to u_1. (Thus, for example, a triangle in G with vertices set $\{a, b, c\}$ gives rise to six 3-cycles abc, acb, bca, bac, cab and cba. Note that, in general, an N-cycle might have repeated vertices.) We now have

$$\sum_i (1-\lambda_i)^N = \mathrm{tr}(B^N) = \sum_C \frac{1}{d_{u_1} d_{u_2} \ldots d_{u_N}} \tag{4.6}$$

the sum being over all N-cycles $C = u_1 u_2 \ldots u_N$ in G. Therefore, $\sum_i (1-\lambda_i)^N$ counts the number of N-cycles, normalised by the degree of each node in the cycle.

The number of N-cycles is related to various graph properties. The number of 2-cycles is just (twice) the number of edges and the number of 3-cycles is (six times) the number of triangles. Hence $\sum_i (1-\lambda)^3$ is related to the clustering coefficient, as discussed below. An important graph property is the number of 4-cycles. A graph which has the minimum number of 4-cycles, for a graph of its density, is quasi-random, i.e., it shares many of the properties of random graphs, including, typically, high connectivity, low diameter, having edges distributed uniformly through the graph, and so on. This statement is made precise in [Tho87] and [CGW89]. For regular graphs, (4.6) shows that the sum $\sum_i (1-\lambda)^4$ is directly to the number of 4-cycles. In general, the sum counts the 4-cycles with weights: for the relationship between the sum and the quasi-randomness of the graph in the non-regular case, see the more detailed discussion in [Chu97, Chapter 5]. The right hand side of (4.6) can also be seen in terms of random walks. A random walk starting at a vertex with degree d_u will choose an edge with probability $1/d_u$ and at the next vertex, say v, choose an edge with probability $1/d_v$ and so on. Thus the probability of starting and ending randomly at a vertex after N steps is the sum of the probabilities of all N-cycles that start and end at that vertex. In other words exactly the right hand side of (4.6). As pointed out in [WL06], random walks are an intricate part of the Internet AS structure.

The left hand side of Equation (4.6) provides an interesting insight into graph structure. The right hand side is the sum of normalised N-cycles whereas the left hand side involves the spectral decomposition. We note in particular that the spectral gap is diminished because eigenvalues close to one are given a very low weighting compared to eigenvalues far from one. This is important as the eigenvalues in the spectral gap typically represent “random” links in the network and are not therefore important parts of the larger structure of the network.

Next, I consider the well-known clustering coefficient. It should be noted that there is little connection between the clustering coefficient, and cluster identification, referred to above. The clustering

[3]Trace of a square matrix is the sum of the elements in the main diagonal

coefficient, $\gamma(G)$, is defined as the average number of triangles divided by the total number of possible triangles

$$\gamma(G) = 1/n \sum_i \frac{T_i}{d_i(d_i-1)/2}, d_i \geq 2 \tag{4.7}$$

where T_i is the number of triangles for node i and d_i is the degree of node i. Now consider a specific triangle between nodes a, b and c. For the cluster coefficient, noting that the triangle will be considered three times, once from each node, the contribution to the average is

$$\frac{1}{d_a(d_a-1)/2} + \frac{1}{d_b(d_b-1)/2} + \frac{1}{d_c(d_c-1)/2} \tag{4.8}$$

However, for the weighted spectrum (with $N = 3$), this particular triangle gives rise to six 3-cycles and contributes

$$\frac{6}{d_a d_b d_c} \tag{4.9}$$

So, it can be seen that the clustering coefficient normalises each triangle according to the total number of possible triangles while the weighted spectrum (with $N = 3$) instead normalises using a product of the degrees. Thus, the two metrics can be considered to be similar but not equal. Indeed, it should be noted that the clustering coefficient is in fact not a metric in the strict sense. While two networks can have the same clustering coefficient they may differ significantly in structure. In contrast, the elements of $\sum_i (1-\lambda)^3$ will only agree if two networks are isomorphic.

The *weighted spectrum* is formally defined as the normalised sum of N-cycles as

$$W(G,N) = \sum_i (1-\lambda_i)^N \tag{4.10}$$

However, calculating the eigenvalues of a large (even sparse) matrix is computationally expensive. In addition, the aim here is to represent the *global* structure of a graph and so precise estimates of *all* the eigenvalue values are not required. Thus, the distribution[4] of eigenvalues is sufficient. In this chapter the distribution of eigenvalues $f(\lambda = k)$ is estimated using pivoting and Sylvester's Law of Inertia [Syl52] to compute the number of eigenvalues that fall in a given interval. A measure of the graph can then be constructed by considering the distribution of the eigenvalues as

$$\omega(G,N) = \sum_{k \in K} (1-k)^N f(\lambda = k) \tag{4.11}$$

where the elements of $\omega(G,N)$ form the *weighted spectral distribution*:

$$WSD : G \rightarrow \Re^{|K|}\{k \in K : ((1-k)^N f(\lambda = k))\} \tag{4.12}$$

In addition, a metric can then be constructed from $\omega(G)$ for comparing two graphs, G_1 and G_2, as

$$\Im(G_1, G_2, N) = \sum_{k \in K} (1-k)^N (f_1(\lambda = k) - f_2(\lambda = k))^2 \tag{4.13}$$

where f_1 and f_2 are the eigenvalue distributions of G_1 and G_2 and the distribution of eigenvalues is estimated in the set K of bins $\in [0, 2]$. Equation (4.13) satisfies all the properties of a metric (see Appendix A).

[4]The eigenvalues of a given graph are deterministic and so *distribution* here is not meant in a statistical sense.

Haddadi *et al.* [HFJ+08] consider 3 and 4 to be suitable values of N for the current application: $N = 3$ is related to the well-known and understood clustering co-efficient; and $N = 4$ as a 4-cycle represents two routes (i.e., minimal redundancy) between two nodes. For other applications, other values of N may be of interest.

4.4 Tuning the Topology Models

The aim of this section is to examine how well the topology generators match the Skitter topology for different values of their parameters. To facilitate this comparison, grids are constructed over the possible values of the parameter spaces and various cost functions are evaluated as follows:

1. A cost function measuring the matching between the number of links in skitter and the generated topologies:

$$C_1(\theta) = (l_t(\theta) - l_{skitter})^2 \tag{4.14}$$

   where C_1 is the first cost function, θ are the model parameters (which differ for each topology generator), l_t is the number of links (which is a function of the parameters) and $l_{skitter}$ is the number of links in the Skitter dataset.

2. A cost function measuring the matching between the spectra of the Skitter network and of the generated topologies:

$$C_2(\theta) = \sum_i (P(\Lambda \leq \lambda_{t,i}) - P(\Lambda \leq \lambda_{skitter,i}))^2 \tag{4.15}$$

   where $\lambda_{t,i}$ is the i^{th} eigenvalue for topology t.

3. A cost function measuring the matching of the weighted spectra:

$$C_3(\theta) = \sum_i ((w * P(\Lambda = \lambda_{t,i}) - w * P(\Lambda = \lambda_{skitter,i}))^2 \tag{4.16}$$

   where weight $w = (1 - i)^4$.

The objective of the optimisation is to minimise the sum squared errors between the cost function for skitter and the generated topology. In addition to examining different parameter values across a grid, the optimum parameters with respect to $C_3(\theta)$ are estimated using the Nelder Meade simplex search algorithm [NM65, DW87]. Note that the topologies generated by the topology generators are random in a statistical sense, due to differing random seeds for each run. Ten topologies are generated for each value of θ and the average spectral distribution is calculated. I found that the variance of the spectral distributions was sufficiently low to allow reasonable estimates of the minima in each case.

4.4.1 Link Densities

Figure 4.1 displays the value of the cost function $C_1(\theta)$ as a function of the topology generator parameters. On the upper and lower left graphs, the grayscale colour indicates the value of the cost function.

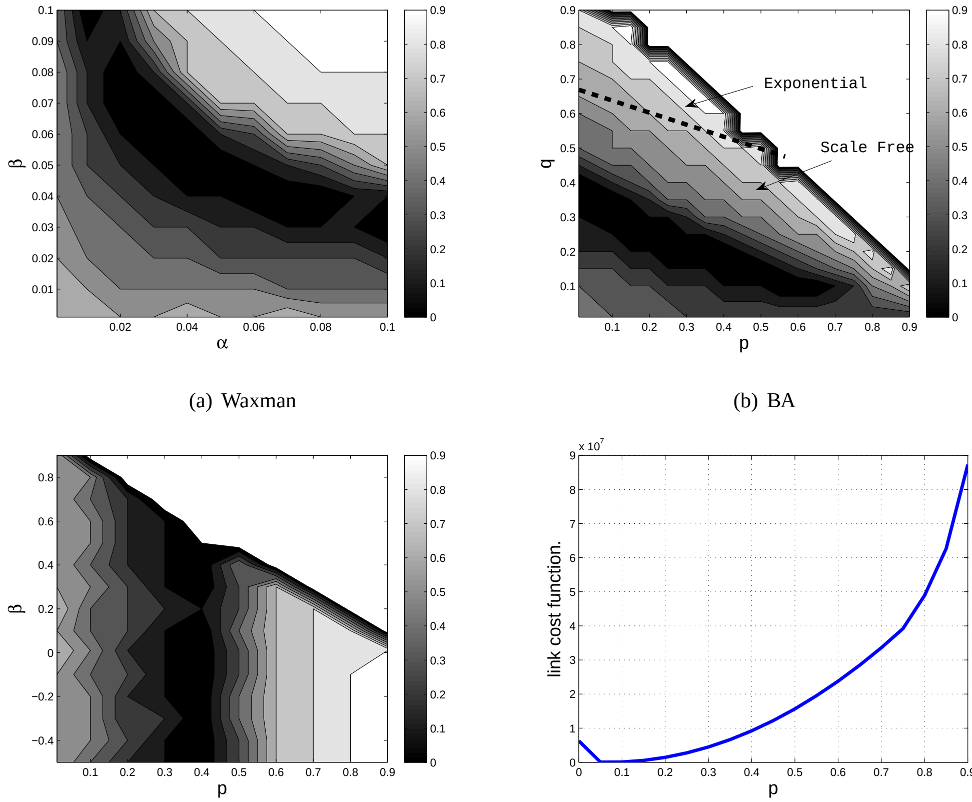

(a) Waxman

(b) BA

(c) GLP

(d) Inet

Figure 4.1: Topology generator parameter grid for sum squared error from number of links.

The darker the region is, the closer the value is to optimal. For Inet (lower right) there is only one parameter, p, so it is plotted as a curve in Figure 4.1(d). Figure 4.1 shows that a minimum exists for each topology in approximately the same regions as the default values of each generator.[5]

For the BA generator it is known that for values of p and q above the line shown in Figure 4.1(b), the topologies generated follow an exponential node degree distribution while those below follow a scale-free distribution. It is encouraging to note that the values of $C_1(\theta)$ are large in the exponential region and the minimum is in the scale-free region as the node degree distribution of the Internet is known to be approximately scale free [AB00]. Overall the results obtained by tuning the parameters based on $C_1(\theta)$ appear reasonable. For link density matching it is possible to obtain parameter values which match the link densities exactly. Indeed, there is a ridge of parameters for BA, GLP and Waxman for which the link densities can be matched. However, as noted in the introduction, there is no control over any other characteristic of the graph using this method.

4.4.2 Spectra PDF

Figure 4.2 shows the spectral PDF of the Skitter dataset and the four topology generators calculated at three parameters values in each grid (the parameter values are indicated in brackets in the legends). The aim is to illustrate how much the spectral PDFs change with the values of the parameters. The spectral

[5] Some of these default values are listed in table 4.1.

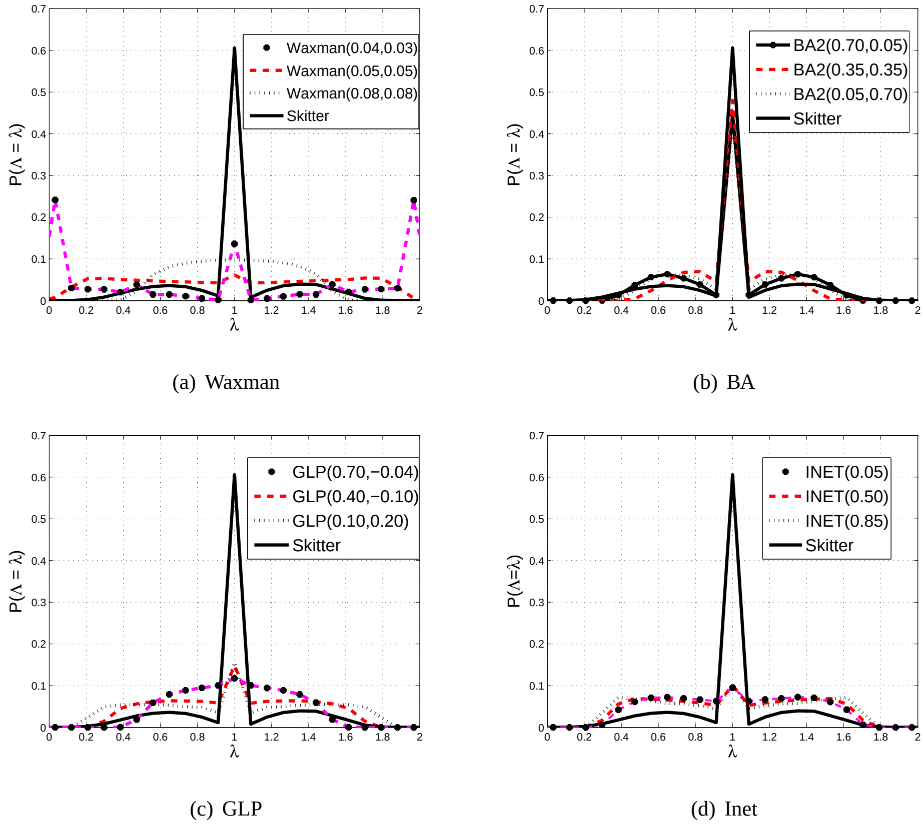

(a) Waxman

(b) BA

(c) GLP

(d) Inet

Figure 4.2: PDF of Spectra

PDFs of Waxman (Figure 4.2(a)) vary significantly for different values of α and β. Furthermore, none of the Waxman PDFs match well the spectral PDF of the Skitter graph. The BA PDFs vary to a lesser extent (Figure 4.2(b)) and appear to give a much better match than the Waxman model, especially around eigenvalue 1 ($\lambda = 1$). This better match of BA is not surprising as the Waxman model is not a good model for the Internet as noted in Section 3.3. GLP (Figure 4.2(c)) and Inet (Figure 4.2(d)) give similar results to BA, with a poor match outside eigenvalue 1. The better match of the BA model around eigenvalue 1 is interesting. As noted in Section 4.3 the regions away from eigenvalue 1 are far more important than the region around $\lambda = 1$. However, what is required is a technique that reveals the differences with distance from one as these are more important. Thus it would appear difficult to evaluate which model, or even which parameter, is better based on the PDFs alone. This point is now further explored by analysis of the grids calculated with respect to $C_2(\theta)$.

4.4.3 Limitations of Spectra CDF

Figure 4.3 shows the value of the second cost function $C_2(\theta)$ as a function of the topology generator parameters, in the same way as Figure 4.1. As can be seen in Figure 4.3, there are many islands corresponding to local minima. The variance in the PDFs referred to in this section is actually greater than any gradient that might exist in the grid. This means that it is not possible to estimate the minimum with respect to $C_2(\theta)$. Figure 4.3 shows that the spectrum on its own is not sufficient to identify the optimum parameters of any of the topology generators. This is because each eigenvalue in $C_2(\theta)$ is weighted

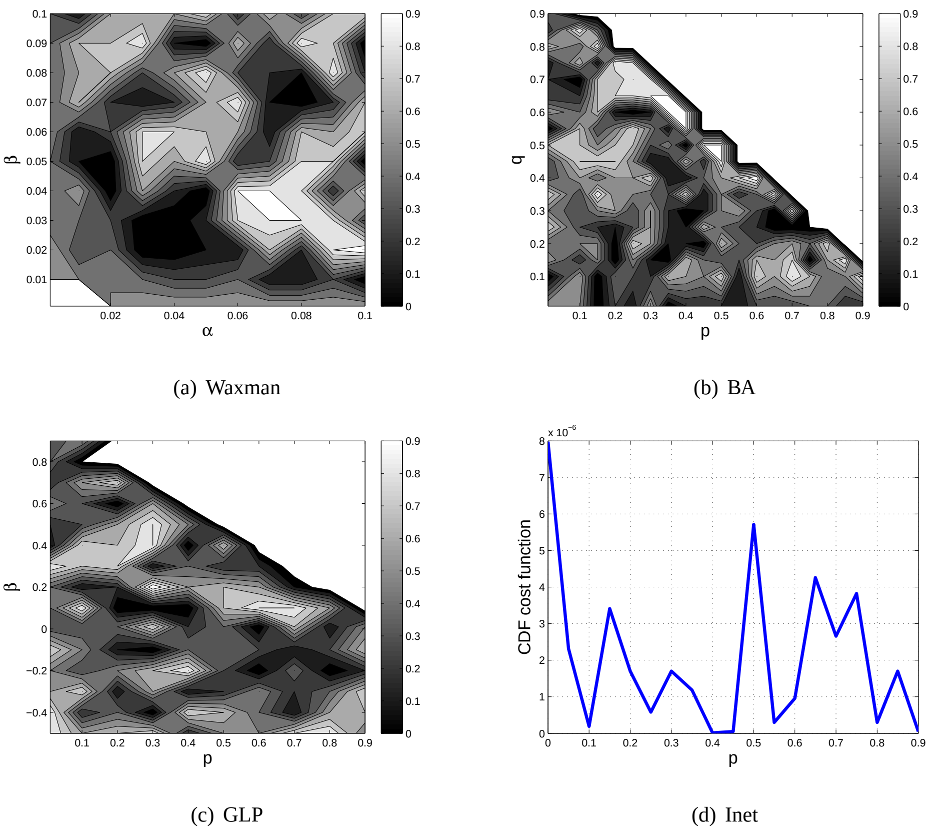

(a) Waxman (b) BA

(c) GLP (d) Inet

Figure 4.3: Parameter grid for sum of absolute differences of spectra CDFs.

equally. As noted in Section 4.3, the eigenvalues close to 1 are more likely to be affected by the random seeds for each topology generator and are the source of the noise on the grid.

4.4.4 Weighted Spectra

The previous section illustrated the limitations of using the raw eigenvalues to find optimal topology generator parameters to match the Skitter topology. Figure 4.4 shows a plot of the weighted spectra of the same topologies as those shown on Figure 4.2. As it can be seen the results are quite different from those shown in Figure 4.2. The Waxman weighted spectra still shows a bad fit with respect to the Skitter data (mainly around 0 and 2) compared to the other generators. The other generators (BA, GLP and Inet) now show that they are capable of matching the weighted spectra of the Skitter topology, especially around the point of greatest weight ($\lambda = 0.4$ or 1.6). The difference between the weighted spectra around 1 is no longer of importance (in contrast to Figure 4.2), reflecting that the weights here approach zero as we approach eigenvalue 1. In the next section the optimum values and the resulting weighted spectra will be compared.

4.4.5 Weighted Spectra Comparison

Figure 4.5 shows the grids associated with $C_3(\theta)$. Unlike the spectra in Figure 4.3 where it was difficult to find an optimum minima, the weighting process, hence giving less importance to noisy eigenvalues in the middle and more importance to the significant ones, has made it possible to get get an optimum region for the parameters. As can be seen the grids show that there is a region with a minima in each case

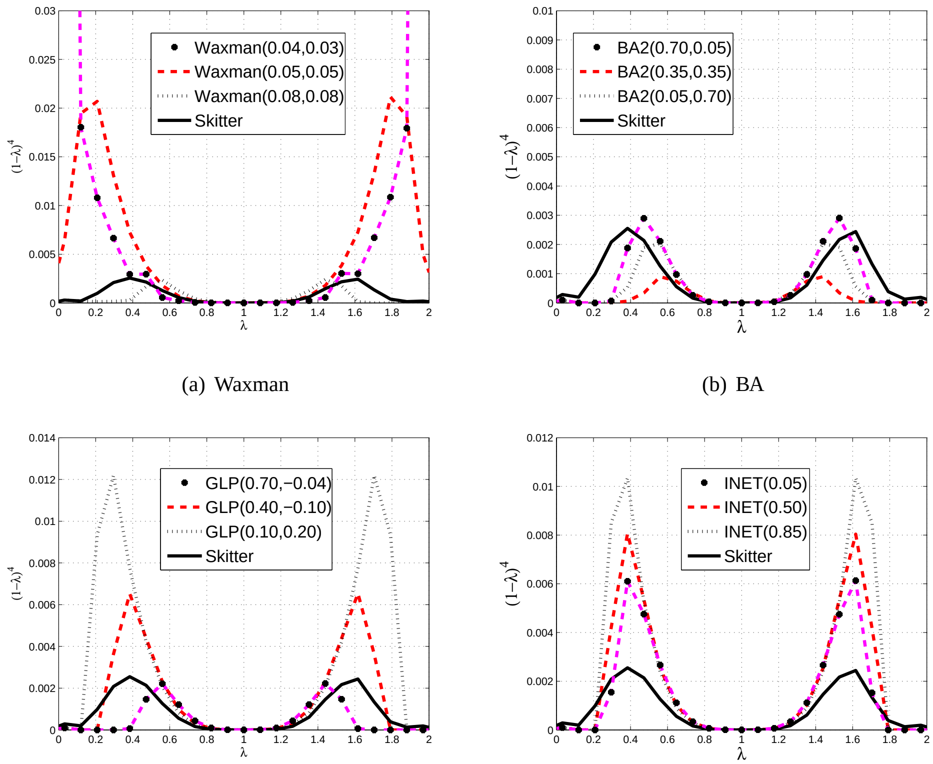

(a) Waxman (b) BA

(c) GLP (d) Inet

Figure 4.4: Weighted spectra grid for generator parameters.

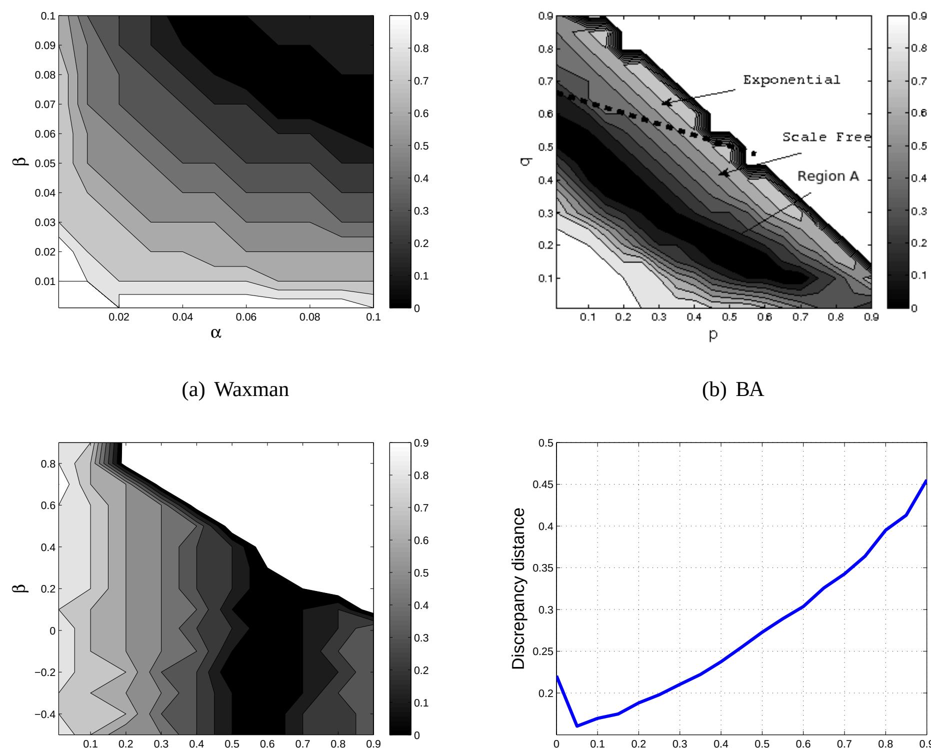

(a) Waxman (b) BA

(c) GLP (d) Inet

Figure 4.5: Grid of sum squared error of weighted spectra for topology generators

Table 4.1: Optimum parameter values for matching Skitter topology.

Waxman	$\alpha = 0.08$ (default$= 0.15$)	$\beta = 0.08$ (default$= -0.2$)	$C_3(\theta) = 0.0026$	$\overline{C_3}(\theta) = 0.0797$
BA	$p = 0.2865$ (default$= 0.6$)	$q = 0.3145$ (default$= 0.3$)	$C_3(\theta) = 0.0014$	$\overline{C_3}(\theta) = 0.0300$
GLP	$p = 0.5972$ (default$= 0.45$)	$\beta = 0.1004$ (default$= 0.64$)	$C_3(\theta) = 0.0021$	$\overline{C_3}(\theta) = 0.0446$
Inet	$\alpha = 0.1013$ (default$= 0.3$)	–	$C_3(\theta) = 0.0064$	$\overline{C_3}(\theta) = 0.0150$
PFP	–	–	$C_3(\theta) = 0.0014$	$\overline{C_3}(\theta) = 0.0371$

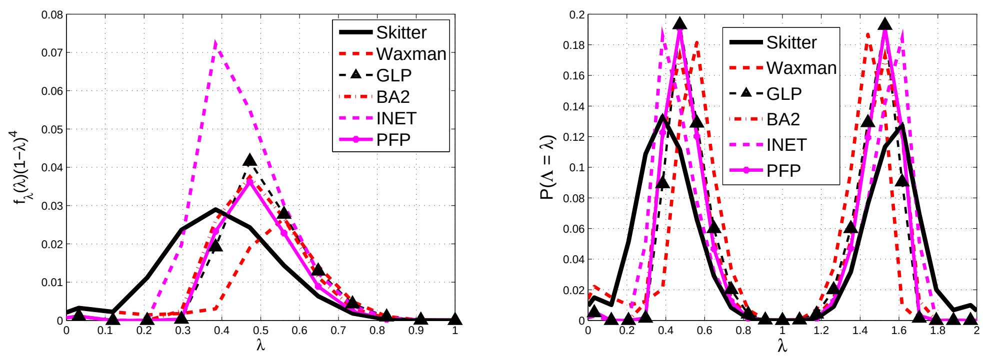

(a) Weighted spectra (b) Normalised weighted spectra

Figure 4.6: Comparison of the weighted spectra.

and in addition, comparing Figure 4.5 and Figure 4.1 it can be seen that these minima lie in a region close to those for $C_1(\theta)$. However, it should be noted that the weighted spectra will try to fit more than just the number of links in a topology. This demonstrates the inherent trade-off. Also of note is that the region of interest for the BA model lies inside the region of scale-free behaviour as shown in Figure 4.5(b).

4.5 Generating Topologies with the Optimum Value Parameters

Table 4.1 displays the optimum values for the topology generators for generating networks that are close to the Skitter graph. In addition, I give the values for $C_3(\theta)$, which show that PFP gives the closest fit followed by BA, GLP, Waxman and finally Inet. While these results are mostly expected, the ranking of Inet as the worst topology generator is surprising. I have also listed some of the default parameters used in certain generators such as BRITE [MLMB01]. While many of the optimised parameters are close to the default values, which is encouraging, it should be noted that the default parameters given by designers are for a *typical* graph and are not selected for any particular situation (e.g., Skitter in this example). Thus a direct comparison is meaningless and it can be seen that optimum parameters are sometimes significantly different from the default ones.

Figure 4.6(a) shows the weighted spectra for each of the topology generators and inspection of this figure goes some way to explaining the discrepancy in the results. As can be seen the main peak in the weighted spectra for the Skitter data occurs at a value of $\lambda = 0.4$. The Waxman generator peak occurs at λ = 0.6 which is closer to 1 demonstrating the greater amount of random structure in the Waxman topologies. However, for the Inet generator the peak occurs at the correct point (λ =0.4) but the weighted power at this point is far greater than in the skitter topology. By normalising the weighted spectrum this

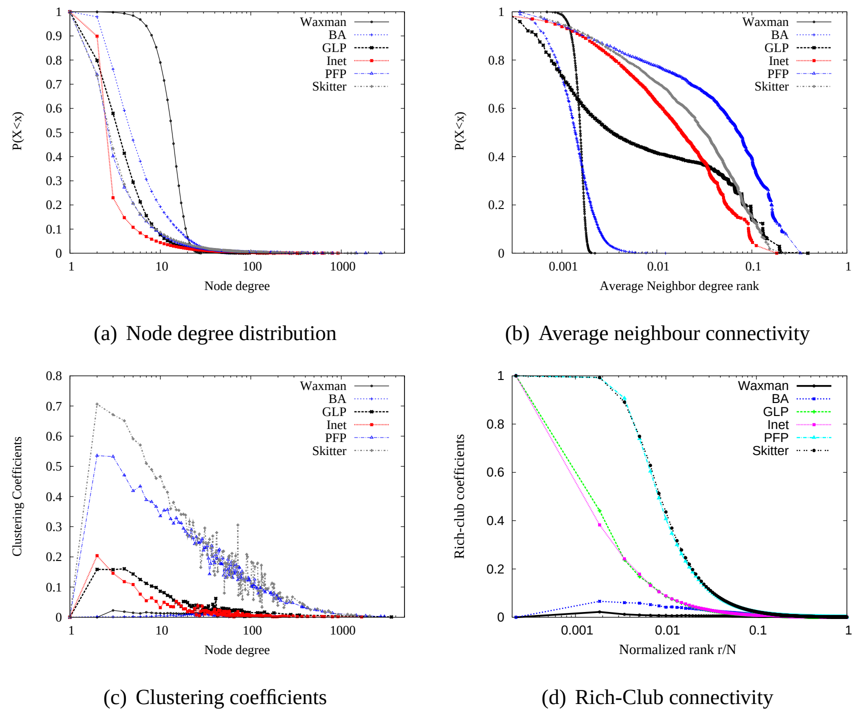

(a) Node degree distribution

(b) Average neighbour connectivity

(c) Clustering coefficients

(d) Rich-Club connectivity

Figure 4.7: Comparison of topology generators and Skitter topology.

point becomes clear:

$$\overline{C_3}(\theta) = \sum_i \frac{((w_i * P(\Lambda = \lambda_{t,i}))}{\sum_i((w_i * P(\Lambda = \lambda_{t,i}))} - \frac{((w_i * P(\Lambda = \lambda_{skitter}))}{\sum_i((w_i * P(\Lambda = \lambda_{skitter}))} \tag{4.17}$$

Using the normalised weighted spectrum the results in Figure 4.6(b) show that Inet is the best match for the Skitter data while the Waxman model still performs worse than the other models. Further research is required before stating which version of C_3 is superior.

Figure 4.7 shows a comparison of the optimised topologies with respect to four typical network metrics: the node degree distribution, the average neighbour connectivity, the clustering coefficient and the rich-club connectivity [Zho06]. As can be seen PFP gives the best match for these metrics in agreement with the proposed metric $C_3(\theta)$. The performance of the other topologies is mixed showing that while one topology is able to match one metric it fails to match another. For example, the GLP generator achieves a reasonable match for the node degree distribution but fails to match the average neighbour connectivity. This demonstrates that for a weak underlying model (e.g., Waxman) the optimisation can not significantly improve its performance when compared to the Internet AS topology.

4.6 Conclusions and Contributions

Comparison of graph structures is a frequently encountered problem across a number of problem domains. To perform a useful comparison requires definition of a cost function that encodes which features of the graphs are considered important. Although the spectrum of a graph is often claimed to be a way to encode a graph's features, the raw spectrum contains too much noise to be useful on its own. In this chapter I have introduced a new cost function, the *weighted spectral distribution,* that improves on the graph spectrum by discounting those eigenvalues that are believed to be unimportant and emphasising the contribution of those believed to be important.

I use this cost function to optimise the selection of parameter values within the particular problem domain of Internet topology generation. The weighted spectrum was shown to be a useful cost function in that it leads to parameter choices that appear sensible given prior knowledge of the problem domain, i.e., are close to the default values and, in the case of the BA generator, fall within the expected region. In addition, as the metric is formed from a summation, it is possible to go further and identify which particular eigenvalues are responsible for significant differences. Although it is currently difficult to assign specific features to specific eigenvalues, it is hoped that this feature of the cost function will be useful in the future[6].

[6]The work presented in this chapter is the result of collaboration with Damien Fay, Steve Uhlig and my advisors. Damien Fay was mainly responsible for the theory behind weighted spectral distribution. Steve Uhlig contributed to the use of the weighted spectral distribution. However, the largest part of the underlying ideas behind the work and tuning approach, the simulation code and the detailed analysis of the results have been done by me.

Chapter 5

Evolution and Scaling of Internet Topologies

In this chapter I study the evolution of the AS topology as inferred from two different datasets over a period of seven years. To focus on structural changes in the topology, I use the *weighted spectral distribution* as this metric reveals differences in the hierarchical structure of two graphs. The results indicate that the Internet is changing from a core-centred, strongly customer-provider oriented, disassortative network, to a soft-hierarchical, peering-oriented, assortative network. In addition, I use a variety of other metrics to analyse the structural disagreement revealed in the AS topologies inferred from the two different datasets. This disagreement is due to the nature of the measurement techniques. I find that the traceroute dataset has increasing difficulty in sampling the periphery of the AS topology, while the BGP dataset does not sample the inner-most core of the network.

5.1 Introduction

The Internet continuously evolves: new networks are created, old ones disappear, and existing ones grow or merge. At the same time, business dynamics cause interconnections between networks to change. Both these effects cause the underlying topology of the Internet to be in a constant state of flux. Studying the evolution of this topology is important as it impacts a variety of factors relevant to network users and application designers, such as scalability, performance and business incentives. For example, different network structures affect the propagation of both legitimate (e.g., routing) and illegitimate (e.g., viruses) information.

Most efforts to understand the structure of the Internet have focused on the AS topology. There are over 25,000 ASes, each representing a single administrative authority with its own network and peering policies. Thus, the AS topology is a graph reflecting the interconnections between the networks that compose the Internet. Relationships between ASes are typically classified as either customer-provider or peer-peer. Note that as the Internet has grown, many larger networks have come to be represented as more than one AS (i.e., to advertise more than one AS number). As a result, the AS topology may contain edges that do not directly represent a business relationship between two distinct networks. However, the AS topology serves as an available, albeit approximate, measure of the complexity of the Internet's structure at a network level.

Characterising the structure of the AS topology has proved difficult, but it is usually simplified to:

a richly connected core, including the fully meshed tier-1 Internet Service Providers (ISPs), providing connectivity for the huge number of smaller ISPs and customer networks at the periphery of the network. These edge ISPs may connect to only a single upstream provider, or may connect to many for resilience, performance and cost reasons. Recent work has shown that the trend is for networks to try to connect directly in the periphery of the Internet, rather than to the core, bypassing the largest providers [GALM08].

In this chapter I analyse the evolution of the AS topology using two significant datasets, each generated by a different measurement technique: the Skitter dataset using traceroute, and the UCLA dataset using BGP. I focus on the overall structure of the topology, rather than local features such as node degree, using a recently introduced metric called the *weighted spectral distribution* (WSD). This allows us to distinguish topologies with different mixing properties, i.e., how much the core can be differentiated from the periphery of the topology [HFU+08]. A clear distinction between the core and the periphery is believed to be one of the strongest features of the Internet topology [SARK02, Zho06].

This chapter makes three contributions. First, I demonstrate how WSD, as explained in Section 4.3, depicts the mixing between core and periphery in the AS topology in Section 5.3. Second, I find that the AS topology has evolved from a highly hierarchical graph with a clearly distinct core towards a "softer" hierarchy where the core and non-core parts of the topology are less distinct (Section 5.4). Third, I show how the two different measurement techniques, traceroute and BGP, both provide limited but complementary coverage of the AS topology: the traceroute dataset has increasing difficulty sampling the periphery, while the BGP dataset does not sample the Internet's core (Section 5.5).

5.2 Related Work

In this section I outline related work, classified into three groups: evolution of the AS topology, spectral graph analysis of the AS topology, and analysis of the clustering features of the AS topology.

Shyu *et al.* [SLH06] study the evolution of a set of topological metrics computed on a set of observed AS topologies. The authors rely on monthly snapshots extracted from BGP RouteViews from 1999 to 2006. The topological metrics they study are the average degree, average path length, node degree, expansion, resilience, distortion, link value, and the Normalised Laplacian Spectrum. They find that the metrics are not stable over time, except for the Normalised Laplacian Spectrum.

Oliveira et al. [OZZ07] look at the evolution of the AS topology as observed from BGP data. Note that they do not study the evolution of the AS topology structure, only the nodes and links. They propose a model aimed at distinguishing real changes in ASes and AS edges from BGP routing observation artifacts. I use the extended dataset made available by the authors, in addition to 7 years of AS topology data from an alternative measurement method.

Latapy and Magnien [LM08] address the question of studying the relation between the size of a measurement sample and the corresponding topological properties. Based on AS topologies built from IP-level measurements from Skitter for a period from January 2005 to May 2006, they observe an increase in the average degree and the clustering coefficient when a larger dataset is used.

Wang and Loguinov [WL06] propose the Wealth-Based Internet Topology (WIT) model. Interestingly, central to their model is the notion that each AS picks its connections to maximise local random walks. This characteristic of the structure of the AS topology is particularly targeted by the WSD. However, as this model is not publicly available it is not included in our comparisons

The graph spectrum has been used for a variety of purposes in addition to characterisation of Internet topologies, including space comparison [Han07], graph matching [LH01], cluster identification [NJW02] and topology generator tuning [HFU$^+$08]. Gkantsidis *et al.* [GMZ03] perform a comparison of clustering coefficients using the eigenvectors of the k largest eigenvalues of the adjacency matrices of AS topologies. k is chosen to retain the strongest eigenvectors discarding most of the others. Those retained are then shown to represent finer elements of the Internet structure. The rest of the spectrum is considered unimportant, even though other works have shown that the eigenvalues of the adjacency matrix or the normalised Laplacian matrix can be used to accurately represent a topology [But06], and some specific eigenvalues provide a measure of properties such as robustness of a network to failures [JU07].

Vukadinovic *et al.* [VHE02] were the first to investigate the properties of the AS topology based on the normalised Laplacian spectrum. They observe that the normalised Laplacian spectrum can be used to distinguish between synthetic topologies generated by Inet [WJ02] and AS topologies extracted from BGP data. This results indicates that the normalised Laplacian spectrum reveals important structural properties of the AS topology. However, as noted by Haddadi *et al.* [HFU$^+$08], the spectrum *alone* cannot be used directly to compare graphs as it contains too detailed information about the network structure. I expand on this work by demonstrating how appropriate weighting of the eigenvalues can reveal the structural differences between two topologies.

Wool and Sagie [WS04] propose several clustering algorithms to explore the AS topology using just a snapshot of the Skitter data. They focus on identification of the dominant clusters, although their result is sensitive to the parameters chosen such as the minimum cluster size. The technique I use, the WSD, differs in that it focuses on random cycles instead of clusters and does not require any parameter estimation. In addition, I use the k-core decomposition to analyse the core of the Internet AS topology.

Li *et al.* [LCMF08] perform a similar study to the one presented here. In their work they use several different clustering methods to identify the distribution of clustering features throughout a network. Interestingly, their clustering metric gives similar results for the skitter and routeviews (here called UCLA) datasets, while WSD shows differing results reflecting directly the differing sampling characteristics of these two measurement techniques.

5.3 Mixing Properties of Networks

The synthetic topology generator introduced in this section is intended as a strawman tool that can be adjusted to show the effect of different parts of a topology on the resulting WSD. These topologies are generated using a simple model based on the existence of a network core and a periphery, as do most generative models of the Internet. Figure 5.1 shows a small topology of 500 nodes. All M nodes within the graph are first assigned locations using a uniform distribution. Nodes within a circle of diameter D are then defined as the *core* and nodes outside a circle of diameter $D \times (1 - m)$ as the periphery, where $m \leq 1$ is a factor called the mixing factor. Connections are then assigned between the core nodes using

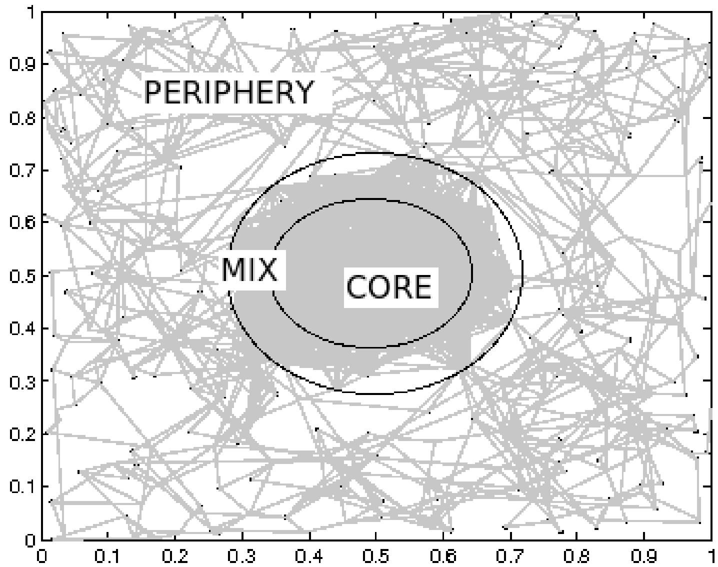

Figure 5.1: Synthetic topology.

a Waxman model:

$$P(u \to v) = \alpha_{core} \exp^{\frac{-d\beta_{core}}{D}} \tag{5.1}$$

where α_{core} and β_{core} are the Waxman coefficients for the core, and d is the distance between two nodes u and v. Subsequently, connections are also assigned in the periphery[1] using a Waxman model but one with different coefficients, α_{per} and β_{per}. After this process, isolated nodes are connected to their nearest neighbour.[2] Figure 5.2 shows the WSD (using $N = 4$) for a topology generated with $M = 2000$ nodes, $D = 0.25$, $\alpha_{core} = 0.08$, $\beta_{core} = 0.08$, $\alpha_{per} = 0.06$, $\beta_{per} = 0.7$, and $m = 0.95$ (i.e., 5% mixing), resulting in a small (relatively) meshed core with a less well connected periphery. There are several things to note in Figure 5.2. Ignoring the asymmetrical part of the curve, which is due to a small number of disconnected components, the peak of the weighted spectrum of the periphery alone lies at $\lambda = 0.7$ while that for the core lies at 0.5. The spectrum for the overall network has *two peaks* at these points. This is a direct result of the fact that the spectrum of a graph is the union of the spectra of its disconnected subgraphs [Chu97]. In terms of the WSD, the union of spectra is equivalent to a weighted average of the WSD. That is, for a graph $G + H$ composed of two disconnected subgraphs G and H:

$$\omega(G + H, N) = |G + H| \left(\frac{\omega(G, N)}{|G|} + \frac{\omega(H, N)}{|H|} \right) \tag{5.2}$$

where $|.|$ denotes volume (number of vertices). Although there is 5% mixing between the core and periphery $\omega(G + H, N)$ results in an close estimate of the network WSD (see Figure 5.2, denoted $\Sigma||E(1 - \lambda_i)^4||$). As $m \to 0$ (i.e., the core and periphery become less and less connected) this estimate becomes more accurate and is exact at m = 0.

[1] Note that nodes lying between D and $D \times (1 - m)$ are members of the core *and* the periphery and will be connected twice.

[2] Note that there are likely to be some disconnected components in the resulting graphs giving asymmetrical spectra, but this does not affect the main results.

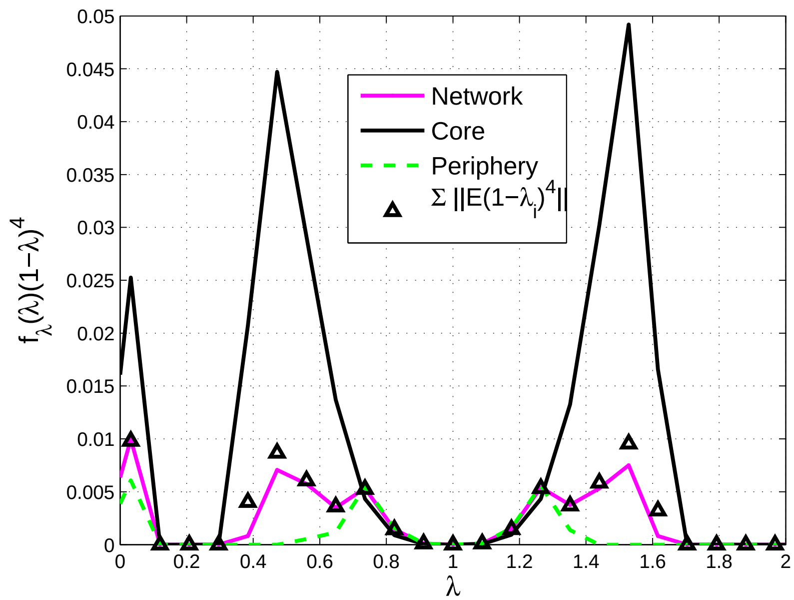

Figure 5.2: Synthetic topology spectra.

Figure 5.3 shows the effect of increasing the mix between the periphery and the core.[3] As can be seen the core becomes less distinct in the resulting spectrum, and has practically disappeared with 40% mixing. Increasing the mixing effectively adds edges connecting the core and periphery, which results in a spreading of the eigenvalues and thus a spreading of the WSD, resulting in less distinct peaks. This result is a consequence of the following theorem from [But07]:

Let G be a weighted graph and H a subgraph on the vertices of G with t non-isolated vertices. If $\{\lambda_0 \leq \lambda_1 ... \leq \lambda_{m-1}\}$ and $\{\theta_0 \leq \theta_1 ... \leq \theta_{n-1}\}$ are the eigenvalues of $L(G)$ and $L(G+H)$ respectively, then for $k = 0, 1, \ldots, n-1$ we have:

$$\lambda_{k+t-1} \leq \theta_k \leq \begin{cases} \lambda_{k-t+1}, & H \text{ is bipartite,} \\ \lambda_{k-t} & \text{otherwise} \end{cases} \tag{5.3}$$

In the current context, the new edges in the mix are being added to t nodes causing the eigenvalues to spread by at most t places. It should be noted that although this makes the core peak less distinct this does not mean that the core is more difficult to detect, rather that the core itself is now less distinct from the periphery.

The statistical properties of the WSD are examined by example in Figure 5.4. This plot was created by generating 50 topologies using the AB model with the optimum parameters using *different initial conditions* and recording the resulting spectra and weighted spectra (as explained in Section 4.4). As the underlying model (i.e. the AB model) is the same for each run, the *structure* might be expected to remain the same and so any structural metric should be relatively robust in the face of varying initial conditions. As can be seen the standard deviation [4] of the (unweighted) spectrum is significantly higher

[3]Again the large peaks before 0.2 represent isolated subgraphs and are ignored.

[4]multiplied by a factor of ten for clarity

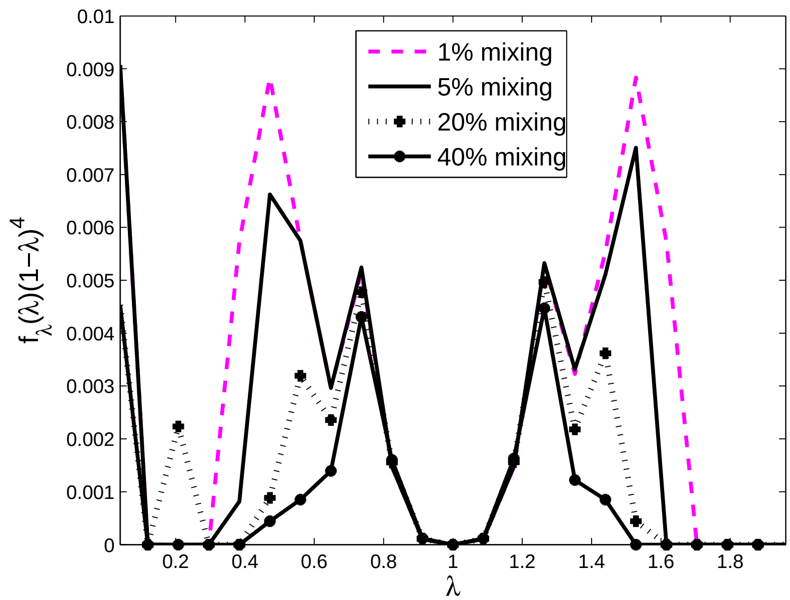

Figure 5.3: Effect of a change in m on the spectrum of the overall network.

at the centre of the spectrum reflecting that the spectral gap contains random connections. However, for the WSD the standard deviation peaks at the same point as the WSD; the noise in the spectral gap having been suppressed.

5.4 Evolution of the Internet

In this section I look at the evolution of the Internet seen through the two datasets using a number of topological metrics. Section 5.4.1 studies the evolution of the AS topology seen in the Skitter dataset, and Section 5.4.2 then studies the evolution of the AS topology seen in the UCLA dataset. We consider the discrepancies between these views in Section 5.5, where I also discuss the likely evolution of the real AS topology.

5.4.1 Skitter topology

The first dataset I study consists of 7 years of traceroute measurements, starting in January 2001, collected by the CAIDA Skitter project [HAA$^+$07]. Traceroutes are initiated from several locations in the world towards a large range of destination IP addresses. The IP addresses reported in the traceroutes are mapped to AS numbers using RouteViews BGP data. I use a monthly union of the set of all unambiguous links collected on a daily basis by the project.[5]

Figure 5.5 presents the evolution over the 7 years of a set of topological metrics computed on the AS topology of Skitter.

The number of ASes seen by Skitter exhibits abrupt changes during the first 40 months. At the end

[5] A link may be ambiguous for a variety of reasons, principally due to problems resolving an IP address to its AS; we ignore such links.

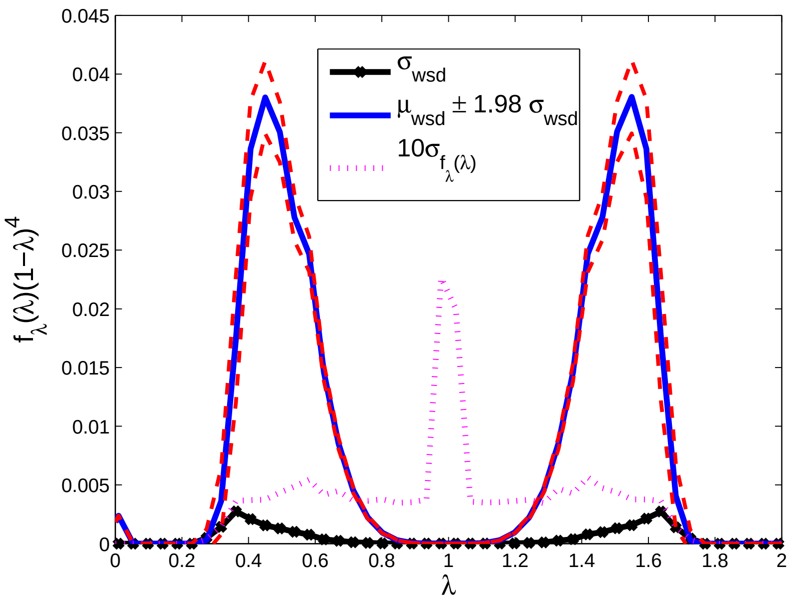

Figure 5.4: Mean and standard deviations for WSD and spectrum for the optimised AB model.

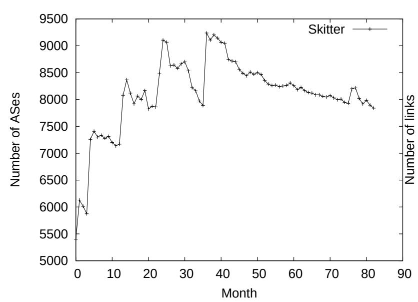

(a) Number of nodes

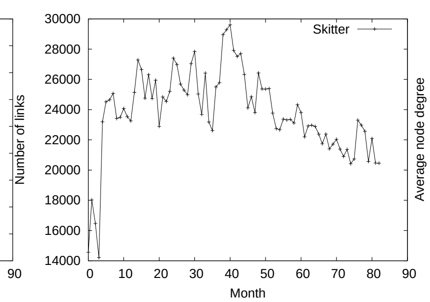

(b) Number of links

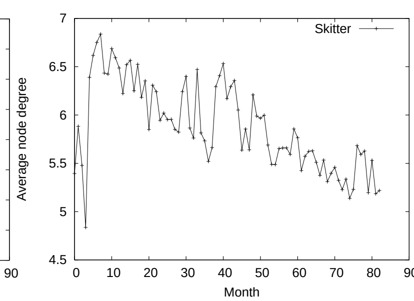

(c) Average degree

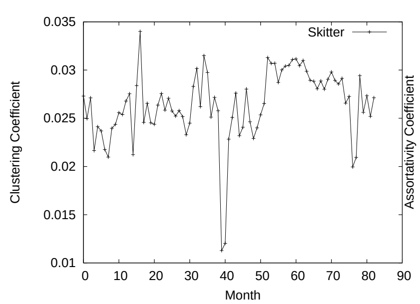

(d) Clustering coefficient

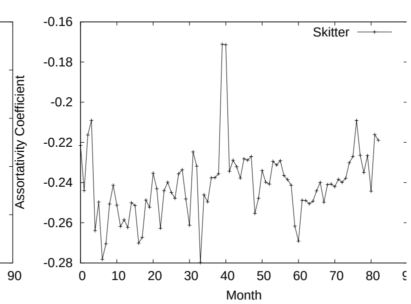

(e) Assortativity coefficient

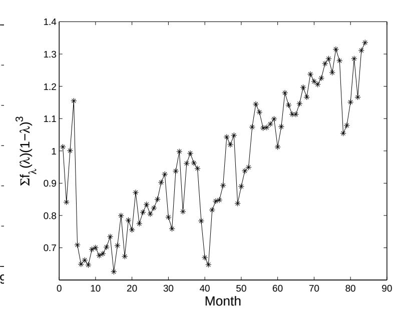

(f) $\omega(G, 3)$

Figure 5.5: Topological metrics for Skitter AS topology.

of those first 40 months, changes were made in the way probing was performed.[6] The large increases in the number of ASes, observed during the first 40 months, are due to new monitors being added to the system. After each increase in the number of ASes a smooth decrease follows, corresponding to a subset of the IP addresses of the Skitter list that no longer respond to probes, e.g., because a firewall starts blocking the probes. The variations in the number of ASes seen by Skitter are not caused by changes in the AS topology itself, but are artifacts of the probing.

The number of AS edges and the average node degree both follow the behaviour of the number of ASes seen. I only observe a large increase in the number of links during the first few months, during which new monitors are added resulting in new regions of the Internet being covered by Skitter measurements. As the list of destinations used by Skitter does not cover the global set of ASes well [BHM$^+$07], and the same list is shared by all monitors, a new monitor will typically discover new ASes close to its location. However, most of the AS edges close to the destination IP addresses have probably already been discovered by existing monitors [LBCX03].

The AS edges that Skitter no longer observes probably still exist but can no longer be seen by Skitter due to its shrinking probing scope. To be effective in observing topology dynamics, traceroute data collection must update destination lists constantly to give optimal AS coverage. This limitation of Skitter is visible in the decreasing average node degree. We would normally expect to see a net increase in the average node degree as ASes tend to add rather than remove peerings, and the results of the BGP data support this view. If the coverage of the Skitter measurements was not worsening, we should see an increasing node degree.

The lower three graphs of Figure 5.5 present the evolution of the clustering coefficient, the assortativity coefficient and the weighted spectrum with $N = 3$, $\omega(G, 3)$ (related to the topology's clustering). We observe that changes were made to the way Skitter probes the Internet around month 40: the metrics take an unusual value, very small for the clustering and very high for assortativity. The values of the clustering and the assortativity coefficients randomly fluctuate over the 7 years, as if the sampling of the AS topology by Skitter is not stable. Neither the clustering nor the assortativity seem to decrease or increase over the 7 years. The value of $\omega(G, 3)$ shows a long-term increasing trend, similar to the decreasing trend in the average node degree. Although related to the clustering, $\omega(G, 3)$ gives different weights to different parts of the topology. The subset of the topology that corresponds to duplicated structures (e.g., the periphery) receives a smaller weight than the rest. The increasing $\omega(G, 3)$ reflects the increasing bias of Skitter toward sampling the core, rather than the periphery, of the Internet.

Figure 5.6 presents four WSDs spanning the entire duration of the Skitter dataset. Notice the eigenvalues at zero, indicating the presence of several disconnected components. The WSD in January 2002 shows a single peak at $\lambda = 0.4$. As time passes, a second peak appears around $\lambda = 0.3$. Thus the sampling resulting in the Skitter data shows an Internet moving from a less hierarchical to more hierarchical topology. This contradicts current observations that AS topology is becoming less hierarchical, with increasing numbers of ASes peering at public Internet Exchange Points (IXPs) to bypass the core

[6] These changes were subject to caveats and bugs affecting measurements, and hence the resulting metrics, at month 40. For more information refer to `http://www.caida.org/data/active/skitter_aslinks_dataset.xml/`

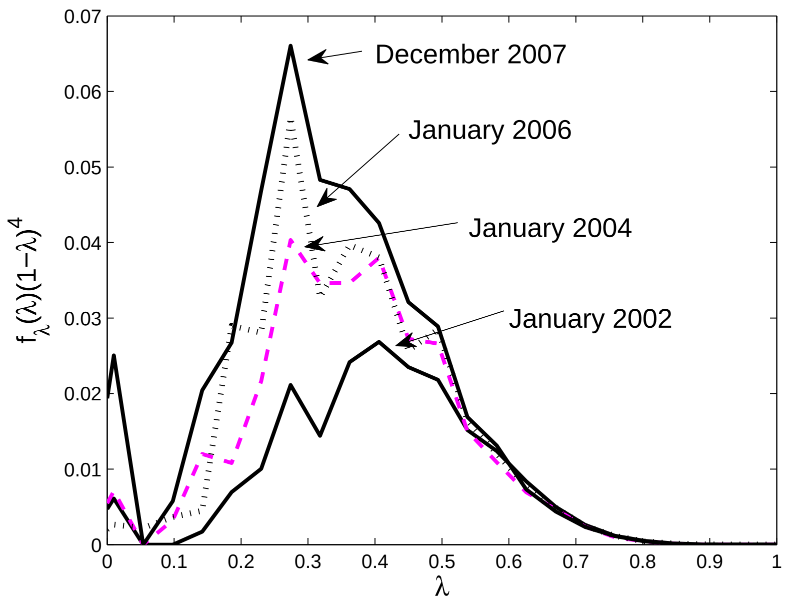

Figure 5.6: Weighted Spectral Distribution, Skitter AS topology.

of the Internet.

To further investigate this surprising result, I next introduce supporting evidence using the k-core measure. A k-core is defined as the maximum connected subgraph, H, of a graph, G, with the property that $d_v \geq k \, \forall v \in H$. As pointed out by Alvarez-Hamelin *et al.* [AHDBV08] the k-core exposes the structure of a graph by pruning nodes with successively higher degrees, k, and examining the maximum remaining subgraph; note this is not the same as simply pruning all nodes with degree k or less. Figure 5.7 shows the proportion of nodes in each k-core as a function of k. There are 84 plots shown but as can be seen there is little difference between each of them demonstrating that the proportion of nodes in each core is not changing over time. This is not surprising due to the nature of the Skitter sampling process: the Skitter data set is composed of traceroutes rooted at a limited set of locations, so the k-core is expected to be similar to *peeling the layers from an onion* [AHDBV08]. From an evolution point of view this result shows that, although the number of nodes being sampled by Skitter is decreasing, the hierarchy of the Internet as observed by Skitter is not changing. This also implies that Skitter is not sampling AS edges and so cannot see evolutionary changes there.

5.4.2 UCLA

I now examine the evolution of the Internet using 52 snapshots, one per month, from January 2004 to April 2008. This dataset, referred to in this chapter as the UCLA dataset, comes from the Internet topology collection[7] maintained by Oliveira *et al.* [OZZ07]. These topologies are updated daily using data sources such as BGP routing tables and updates from RouteViews, RIPE,[8] Abilene[9] and LookingGlass

[7] `http://irl.cs.ucla.edu/topology/`
[8] `http://www.ripe.net/db/irr.html`
[9] `http://abilene.internet2.edu/`

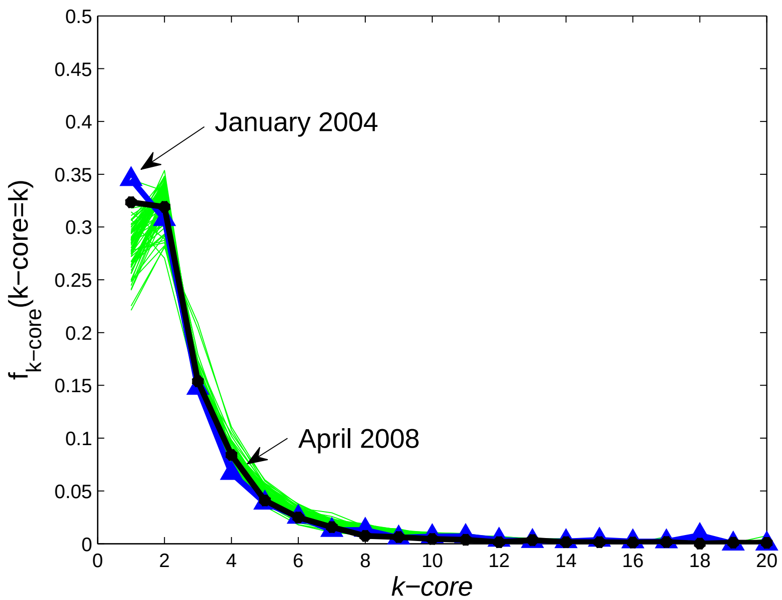

Figure 5.7: k-core proportions, Skitter AS topology

servers. Each node and link is annotated with the times it was first and last observed.

Figure 5.8 presents the evolution of the same set of topological metrics as Figure 5.5, over the 4 years of AS topologies in the UCLA dataset.

The UCLA AS topologies display a completely different evolution to the Skitter dataset, more consistent with expectation. As the three upper graphs of Figure 5.8 show, the number of ASes, AS edges, and the average node degree are all increasing, as expected in a growing Internet.

The increasing assortativity coefficient indicates that ASes increasingly peer with ASes of similar degree. The preferential attachment model is thus becoming less relevant over time. This trend towards a less disassortative network is consistent with more ASes bypassing the tier-1 providers through public IXPs [GALM08], hence connecting with nodes of similar degree. Another explanation for the increasing assortativity is an improvement in the visibility of non-core edges in BGP data. I will demonstrate in Section 5.5 that the sampling of core and non-core edges by UCLA and Skitter biases the observed AS topology structure. Contrary to Skitter, $\omega(G, 3)$ for UCLA decreases over time. As a weighted clustering metric, $\omega(G, 3)$ indicates that the transit part of the AS topology is actually becoming sparser over time compared to the periphery. Increasing local peering with small ASes in order to reduce the traffic sent to providers decreases both the hierarchy induced by strict customer-provider relationships, and in turn decreases the number of 3-cycles on which $\omega(G, 3)$ is based.

If we look closely at Figure 5.9, we see a spectrum with a large peak at $\lambda = 0.3$ in January 2004, suggesting to a strongly hierarchical topology. As time passes, the WSD becomes flatter with a peak at $\lambda = 0.4$, consistent with a mixed topology where core and non-core are not so easily distinguished.

Figure 5.10 shows the proportion of nodes in each k-core as a function of k. There are 52 plots shown as a smooth transition between the first and last plots, emphasised. As can be seen, the distribution

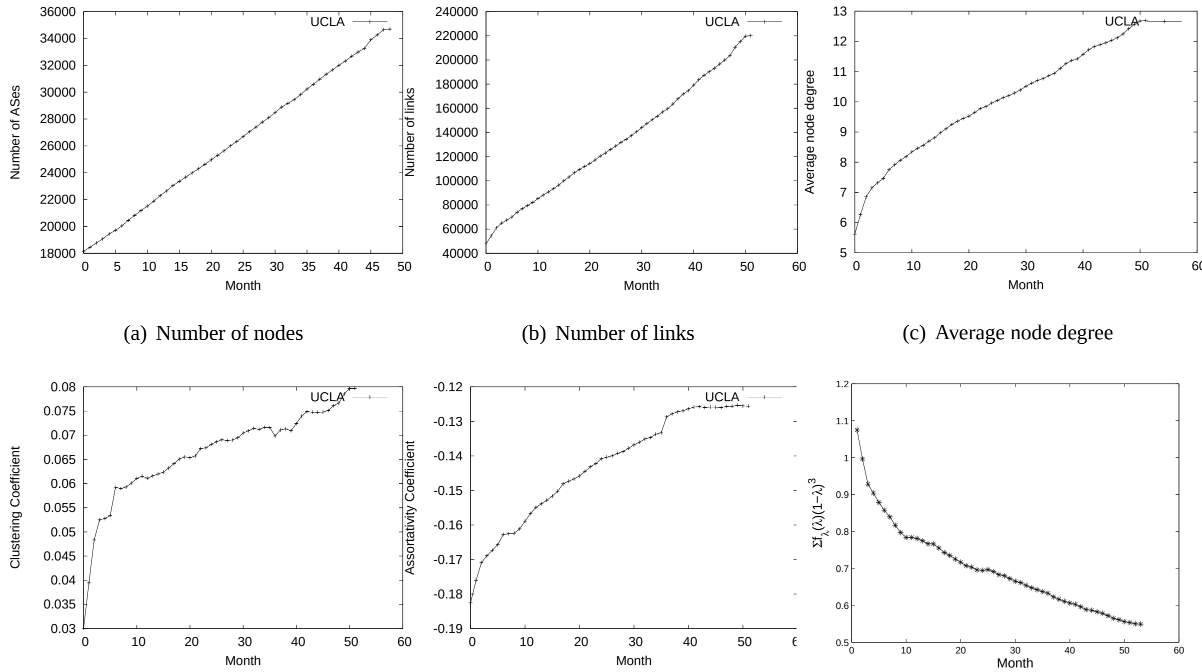

(a) Number of nodes (b) Number of links (c) Average node degree

(d) Clustering coefficient (e) Assortativity coefficient (f) $\omega(G, 3)$

Figure 5.8: Topological metrics for UCLA AS topology.

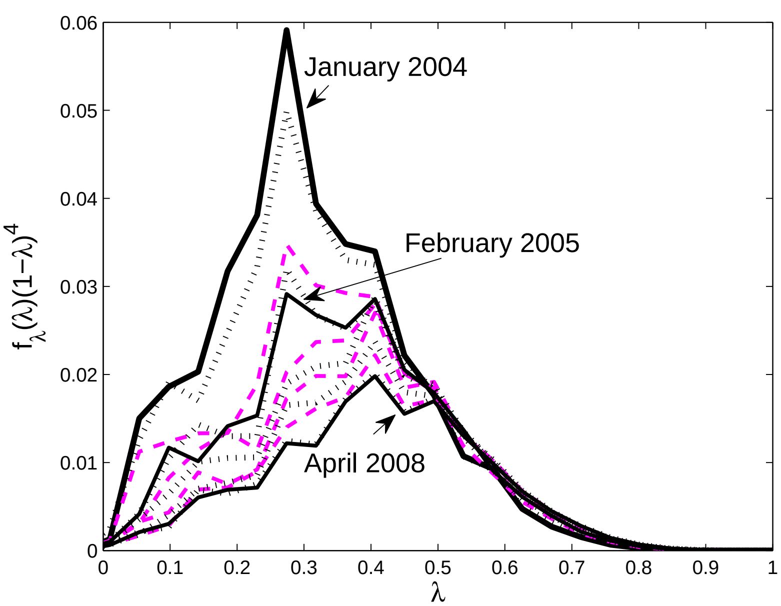

Figure 5.9: Weighted Spectral Distribution, UCLA AS topology.

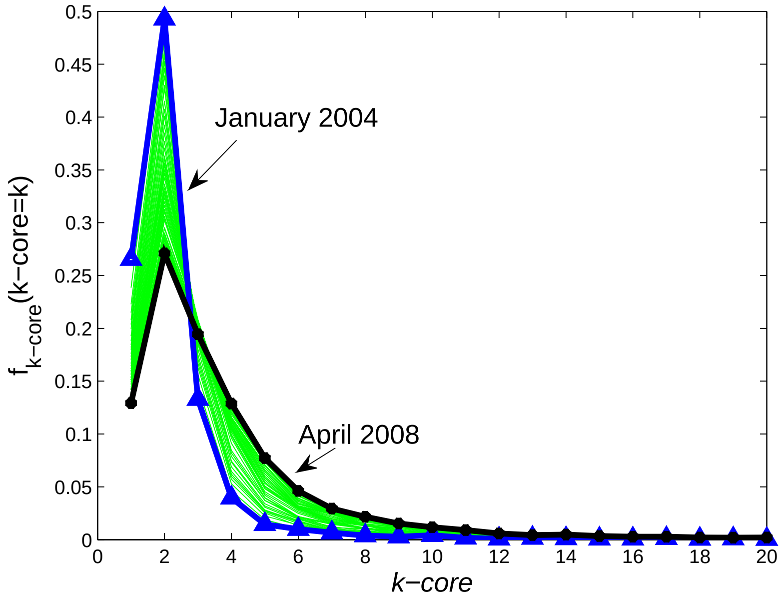

Figure 5.10: k-core proportions, UCLA AS topology.

of k-cores moves to the right over time, indicating that the proportion of nodes with higher connectivity is increasing over time. This adds further weight to the conclusion that the UCLA dataset shows a weakening of hierarchy in the Internet, with more peering connections between nodes. Note that the UCLA data set was not examined in [AHDBV08].

5.5 Reconciling the Datasets

The respective evolutions of the AS topology visible in the Skitter and UCLA datasets differ. Skitter shows an AS topology that is becoming sparser and more hierarchical, while UCLA shows one that is becoming denser and less hierarchical. Can we reconcile those differing views? One must first understand that Skitter and UCLA sample different parts of the AS topology: Skitter sees a far smaller fraction of the real AS topology than UCLA, and even UCLA does not see the whole AS topology [OPW^{+}08].

To check how similar the AS topologies of Skitter and UCLA are, I computed the intersection and the difference between the two datasets in terms of AS edges and ASes. I used a two-years period from January 2006 until December 2007. In Table 5.1 I show the number of AS edges and ASes that Skitter and UCLA have in common during some of these monthly periods (labelled ”intersection”), as well as the number of AS edges and ASes contributed to the total and coming from one of the two datasets only (labelled ”Skitter-only” or ”UCLA-only”). I observe a steady increase in number of total ASes and AS edges seen by the two datasets. At the same time, the intersection between the two datasets decreases. Due to the wide coverage of the UCLA dataset, few ASes and AS edges are contributed by Skitter only.

From Table 5.1, we may conclude that the Skitter dataset is uninteresting. To the contrary, the relatively constant, albeit decreasing, sampling of the Internet core by Skitter gives us a clue about which part of the Internet is responsible for its structural evolution.

	Autonomous Systems				AS Edges			
Time	**Total**	**Intersection**	**Skitter-only**	**UCLA-only**	**Total**	**Intersection**	**Skitter-only**	**UCLA-only**
Jan. 2006	25,301	32.6%	0%	67.4%	114,847	15.4%	5.3%	79.3%
Mar. 2006	26,007	31.6%	0%	68.4%	118,786	14.9%	4.4%	80.7%
May. 2006	26,694	30.5%	0%	69.5%	124,052	13.8%	4.6%	81.5%
Jul. 2006	27,396	29.5%	0%	70.5%	128,624	13.2%	3.7%	83.1%
Sep. 2006	28,108	28.7%	0%	71.3%	133,813	12.6%	3.4%	84.0%
Nov. 2006	28,885	27.9%	0%	72.1%	139,447	12.4%	3.4%	84.2%
Jan. 2007	29,444	27.2%	0%	72.8%	144,721	11.6%	3.1%	85.3%
Mar. 2007	30,236	26.5%	0%	73.5%	151,380	11.2%	3.0%	85.8%
May. 2007	30,978	25.6%	0%	74.4%	157,392	10.5%	2.7%	86.8%
Jul. 2007	31,668	25.9%	0%	86.1%	166,057	10.0%	3.8%	86.2%
Sep. 2007	32,326	24.5%	0%	75.5%	168,876	9.7%	2.5%	87.8%
Nov. 2007	33,001	23.9%	0%	76.1%	174,318	9.5%	2.2%	88.3%

Table 5.1: Statistics on number of ASes and edge counts for datasets

Table 5.2 shows the number of AS edges belonging to the tier-1[10] mesh (labelled ”T1 mesh”) as well as other AS edges where a tier-1 appears. More than 30% of the AS edges sampled by Skitter cross at least a tier-1 AS, against about 15% for UCLA. Both dataset see almost all AS edges from the tier-1 mesh. Note that the decrease in the number of tier-1 edges in Skitter is partly related to IP to AS mapping issues for multi-origin ASes [GALM08].

	Skitter			UCLA		
Time	**Total**	**T1 mesh**	**Other T1**	**Total**	**T1 mesh**	**Other T1**
Jan. 2006	23,805	66	7,498	108,720	64	19,149
Mar. 2006	22,917	66	7,289	113,555	64	19,674
May. 2006	22,888	64	7,504	118,331	64	20,143
Jul. 2006	21,740	65	7,192	123,842	64	20,580
Sep. 2006	21,400	65	6,974	129,228	64	21,059
Nov. 2006	22,034	66	7,159	134,636	65	21,581
Jan. 2007	21,345	65	6,898	140,216	65	22,531
Mar. 2007	21,366	65	6,774	147,000	65	23,194
May. 2007	20,738	65	6,694	153,156	65	23,769
Jul. 2007	22,972	65	6,838	159,792	65	24,310
Sep. 2007	20,570	64	6,510	164,770	65	24,888
Nov. 2007	20,466	64	6,430	170,431	65	25,480

Table 5.2: Coverage of tier-1 edges by Skitter and UCLA.

[10]I rely on the currently accepted list of 12 tier-1 ASes that provide transit-only service: AS174, AS209, AS701, AS1239, AS1668, AS2914, AS3356, AS3549, AS3561, AS5511, AS6461, and AS7018.

The evolution of the AS topology observed by the Skitter and UCLA datasets is not inconsistent as it first appeared from Section 5.4. Rather, the two datasets sample differently the AS topology, leading to different bias. A large fraction of the AS topology sampled by Skitter relates to the core, i.e., edges containing at least a tier-1 AS. With its wider coverage, UCLA observes a different evolution of the AS topology, with a non-core part that grows more than the core. The evolution seen from the UCLA dataset seems more likely to reflect the evolution of the periphery of the AS topology. The non-core part of the Internet is growing and is becoming less and less hierarchical. Despite a common trend towards making a *union* of the datasets in this field, such simple addition is not appropriate for the UCLA and Skitter datasets. Each dataset has its own biases and measurement artifacts. Mixing them together will only add these biases together, potentially leading to poorer quality data. Further research is required in order to devise a correct methodology that takes advantage of different datasets obtained from different sampling processes.

The above observations suggests that the Internet, once seen as a tree-like, disassortative network with strict power law properties [FFF99], is moving towards an assortative and highly inter-connected network. Tier-1 providers have always been well connected, but the biggest shift is seen at the Internet's periphery where content providers and small ISPs are aggressively adding peering links among themselves using IXPs to avoid paying transit charges to tier-1 providers. However, a different view of the Internet evolution can be obtained using the WSD, shown in Figures 5.6 and 5.9. As seen in Section 5.3, one possible cause for this behaviour is increased mixing of the core and periphery of the network, i.e. the strict tiered hierarchy is becoming less important in the network structure. This is given further weight by studies such as [OPW+08] which show that the level of peering between ASes in the Internet has greatly increased during this period, leading to a less core-dominated network.

5.6 Conclusions and Contributions

In this chapter I presented a study of two views of the evolving Internet AS topology, one inferred from traceroute data and the other from BGP data. I exposed discrepancies between these two inferred AS topologies and their evolution. I reconciled these discrepancies by showing that the topologies are not directly comparable as neither method sees the entire Internet topology: BGP data misses some peerings in the core which traceroute observes; traceroute misses many more peerings than BGP in the periphery. However, traceroute and BGP data do provide complementary views of the AS topology.

To remedy the problems of decreasing coverage by the Skitter traceroute infrastructure and the lack of visibility of the core by UCLA BGP data, significant improvements in fidelity could be achieved with changes to the existing measurement systems. The quality of data then collected by the traceroute infrastructure would benefit from greater AS coverage, while the BGP data would benefit from data showing intra-core connectivity. I acknowledge the challenges inherent in these improvements but emphasise that, without such changes, the study of the AS topology will forever be subject to the vagaries of imperfect and flawed data. Availability of traceroute data from a larger number of vantage points, as attempted by the Dimes project, will help remedy these issues. However even such measurements have to be done on a very large scale, and ideally performed both from the core of the network (like Skitter), as well as the edge (like Dimes).

To provide an objective analysis of the changing structure of the AS topology, I used a wide range topological metrics, including the weighted spectral distribution. I find that the core of the Internet is becoming less dominant over time, and that edges at the periphery are growing instead. The practice of content providers and content distribution networks seeking connectivity to greater numbers of ISPs at the periphery, and the rise of multi-homing, both support these observations. Further, I observe a move away from a preferential attachment, tree-like disassortative network, toward a network that is flatter, highly-interconnected, and assortative. These findings are also indicative of the need for more detailed and timely measurements of the Internet topology, in order to build up on works such as [Eco05], focusing on the economics of the structural changes such as institutional mergers, dual homing and increasing peering relationships.[11]

[11] The work presented in this chapter is the result of collaboration with Damien Fay, Andrew G. Thomason, Steve Uhlig and my advisors. Damien Fay and Andrew G. Thomason were mainly responsible for the theory behind weighted spectral distribution. Steve Uhlig contributed to the use of the weighted spectral distribution and helped with understanding of the Internet evolution. However, the largest part of the underlying ideas behind the work and the ideal approach, analysis code, collection and preparation of the traces and the detailed analysis of the results have been done by me.

Chapter 6

Contributions and Future Work

This chapter concludes this thesis by summarising the work carried out, the contributions and suggesting areas of future work.

6.1 Conclusions and Contributions

My main contributions include analysis of popular AS topology generators, comparing them with numerous observations, and highlighting appropriate metrics for comparing the models through long term observations of the evolution of the Internet. The conclusions and contributions can be broken down into categories listed in this section.

6.1.1 Identifying Modelling Challenges

In Chapter 3, I provided insight into the Internet AS topology. I evaluated various models for generating synthetic topologies and compared them to observed AS topologies collected at different times using different measurement methods. I based this comparison on numerous topological and statistical measures.

My analysis revealed that current topology models do not accurately represent the observed Internet AS topology. Although current models accurately preserve the degree-related properties and preferential attachment, they fail to reproduce local connectivity metrics. At the same time, I observe that more recent topology generators generally perform better than older ones. This is partly due to the availability of better observed topologies. I believe that, in addition to degree-related, clustering and centrality properties, the highly meshed core of the Internet AS topology must be considered in order to generate representative synthetic topologies.

I also compared the properties of AS topologies relying on different sets of observations. It was observed that, in contrast to structural metrics, node degree-related properties are not greatly affected by the addition of more vantage points as they add only a small percentage of peering links. On the other hand, the power-law nature of the node degree distribution seems questionable, as increasing the number of observation points causes deviation from strict power-law scaling.

6.1.2 Tuning Topology Generators

A new cost function, the weighted spectral distribution (WSD), was introduced in Chapter 4. The WSD improves on the graph spectrum by discounting those eigenvalues that are believed to be unimportant and emphasising the contribution of those believed to be important.

I used this cost function to optimise the selection of parameter values within the particular problem domain of Internet topology generation. Optimal parameters relative to this cost function were then estimated for Internet topology generators. The WSD was shown to be a useful cost function in that it leads to parameter choices that appear sensible given prior knowledge of the problem domain. It capture wells the clustering characteristics and it is sensitive to mixing between the core and edge ASes. In addition, as the metric is formed from a summation, it is possible to go further and identify which particular eigenvalues are responsible for significant differences. Due to high computational cost of calculating the eigenvalues, it is currently difficult to assign specific features to specific eigenvalues, it is hoped that this feature of the cost function will be useful in the future.

6.1.3 Analysis of the Internet Evolution

In Chapter 5 I presented a study of two views of the evolving Internet AS topology, one inferred from traceroute data and the other from BGP data. I exposed inconsistencies between these two inferred AS topologies and their evolution. I reconciled these inconsistencies by showing that the topologies are not directly comparable as neither method sees the entire Internet topology: BGP data misses many peerings in the core which traceroute observes; traceroute misses many more peerings than BGP in the periphery. However, traceroute and BGP data complement each other.

To remedy the problems of decreasing coverage by Skitter traceroute infrastructure and lack of visibility of the core by UCLA BGP data, significant improvements in fidelity could be achieved with changes to the existing measurement systems. The quality of data then collected by the traceroute infrastructure would benefit from greater AS coverage, while the BGP data would benefit from data showing intra-core connectivity. I acknowledge the challenges inherent in these improvements but emphasise that, without such changes, the study of the AS topology will forever be subject to the vagaries of imperfect and flawed data.

To provide an objective analysis of the changing structure of the topology, I used a wide range of topological metrics, including the WSD. I observed that the core of the Internet is becoming less dominant over time, and that edges at the periphery are growing instead. The practice of content providers and content distribution networks seeking connectivity to greater numbers of ISPs at the periphery, and the rise of multi-homing, both support this hypothesis. Further, I observe a move away from a preferential attachment, tree-like disassortative network, towards a network that is flat, highly-interconnected, and assortative

6.2 Discussions and Future Work

Valuable future work in this area is to consider the analysis for router-level topologies. Such an analysis of router-level topologies is bound to differ greatly from AS-level ones, as network operators have tight control over router interconnects and are subject to different constraints from the AS-level connectivity. The control plane at the router level has different characteristics to those seen at the AS level. At the router level, the dynamics are more frequent and tend to have a shorter durations. Regular maintenance works, router and link failures, traffic engineering, firewall misbehaviours and other factors all effect the routing at the IP layer. Operators do not disclose information about routing changes and link failures. This has made it difficult to model the behaviour of router level Internet topology. I am currently inferring the characteristics of router level topologies of a major tier-1 ISP, looking at short term and long terms trends, while considering the effects of failures on the ISP network topology. This will also allow researchers to build a model for dynamic topology generation at the router level.

Today, topology generators are tightly bound to the observed data used to validate them. Given that the actual properties of the Internet topology are not known, topology generators should strive to reproduce the variability that characterises the evolution of the Internet topology over time. Future topology generators should be able to express the variations in local connectivity that makes today's Internet: peering relationships, internal AS topology and routing policies each changing over time due to failures, maintenance, upgrades and business strategies of the network. Topology generators should capture those dimensions, by allowing a certain level of randomness in the outcome, rather than enforcing structural assumptions as the truths about Internet's evolving structure, which may never be discovered. If incorrect AS interconnections or policies are used for simulation purposes, then the resulting routes might be far from realistic [MFM+06, MUF+07].

The Internet is not a static network. At the AS level, there is a constant growth in the number of peering links between ISPs [OZZ07]. Also, due to policy routing and hot potato routing, the changes at the IP level affect the AS level [TSGR04]. Simulation for applications such as routing protocols and analysis such as studies in prefix hijacking would benefit from topologies which take into account the changes of the network over time, similar to real network behaviour. I believe that using static topologies does not fully exploit the potential scenarios that one should consider in simulations. Another important aspect of the networks that is not captured by current models is the move of the Internet AS topology towards having a meshed core of tier-1 ISPs, alongside multiple peering relationships between edge ASes, and an atypical connection models of some ASes such as the content providers which form many peering connections with as many ASes as possible in order to avoid high transit charges [OPW+08]. In addition to information about peering links, the availability of models of growth and evolution of networks will enable us to include dynamic models for generating synthetic AS topologies. Pursuing this goal, I aim to form a collaboration with network operators, alongside topology generator designers, to provide a representative dynamic topology generator to the research community.

Finally, the metric used in chapter 5, WSD, can be used for analysis of a wide range of topologies and it is not necessarily bound to Internet topologies. As such, I am looking at using the WSD for exploiting the hierarchy and structural characteristics of social networks and protein-protein interaction networks. If successful, this can be a very efficient yet accurate method of categorising such large networks which may be formed of millions of links.

Appendix A

WSD Metric Proof

The WSD metric proposed in Chapter 4 for obtaining a best fit is:

$$J(G_x, G_y) = \sum_{k \in K} (1-k)^4 (f_x(\lambda = k) - f_y(\lambda = k))^2 \tag{A.1}$$

We now show that $\sqrt{J(G_x, G_y)}$ is a metric in the mathematical sense. The difference between $\sqrt{J(G_x, G_y)}$ and $J(G_x, G_y)$ is similar to the difference between the sum squared error and the root mean squared error. We prefer the sum squared error (i.e., $J(G_x, G_y)$) in this application as it provides the well known minimum variance-bias trade-off.

A metric satisfies the following four conditions:

(*a*) $J(G_x, G_y) \geq 0$ (non-negativity)

(*b*) $J(G_x, G_y) = 0 \Leftrightarrow x = y$ (identity of indiscernibles)

(*c*) $J(G_x, G_y) = J(G_y, G_x)$ (symmetry)

(*d*) $J(G_x, G_z) \leq J(G_x, G_y) + J(G_y, G_z)$ (triangle inequality)

(*a*) and (*c*) follow directly from (A.1). Noting that all the elements of the sum in $J(G_x, G_y)$ are positive $\implies J(G_x, G_y) = 0$ if and only if $f_x(\lambda = k) = f_y(\lambda = k)\ \forall k$. Arranging (and increasing the number of bins if necessary) the k bins such that each bin contains at most 1 eigenvalue Requires G_x to be co-spectral and isomorphic to G_y. Two graphs may be *co-spectral*, i.e., they share the same spectrum but are not isomorphic. However, studies have shown [ZW05] that the number of co-spectral graphs falls dramatically with the number of vertices in the graph. For example, only 0.05% of all graphs with 21 vertices are co-spectral and not isomorphic; this number is thought to decrease with increasing number of vertices [ZW05]. Thus, condition (*b*) is true *almost certainly*, in the statistical sense.

$\sqrt{J(G_x, G_y)}$ defines the standard metric space R_2^K [KF75]. This can be seen by distributing the weights $(1-k)^4$ as:

$$\sqrt{J(G_x, G_y)} = \left(\sum_{k \in K} (h_x(\lambda = k) - h_y(\lambda = k))^2 \right)^{1/2} \tag{A.2}$$

where

$$h_x(\lambda = k) = (1-k)^2 f_x(\lambda = k) \tag{A.3}$$

and $h_y(\lambda = k)$ is similarly defined. The triangle inequality holds for (A.2). For a detailed proof see [KF75] Chapter 2, Section 5.

Bibliography

[AB00] Reka Albert and Albert-Laszlo Barabasi. Topology of evolving networks: local events and universality. *Physical Review Letters*, 85:5234, (2000). 44, 76

[ACKM05] Dimitris Achlioptas, Aaron Clauset, David Kempe, and Cristopher Moore. On the bias of traceroute sampling: or, power-law degree distributions in regular graphs. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 694–703, Baltimore, MD, May 2005. 29

[ACL00] William Aiello, Fan Chung, and Linyuan Lu. A random graph model for massive graphs. In *STOC'00: Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pages 171–180, Portland, OR, May 2000. 37, 50, 70

[ACO$^+$06] Brice Augustin, Xavier Cuvellier, Benjamin Orgogozo, Fabien Viger, Timur Friedman, Matthieu Latapy, Clemence Magnien, and Renata Teixeira. Avoiding traceroute anomalies with paris traceroute. In *Proceedings of ACM/Usenix Internet Measurement Conference (IMC) 2006*, pages 153–158, Rio de Janeiro, Brazil, October 2006. 13, 30

[ADGW03] David Alderson, John Doyle, Ramesh Govindan, and Walter Willinger. Toward an optimization-driven framework for designing and generating realistic Internet topologies. *SIGCOMM Computer Communications Review*, 33(1):41–46, 2003. 41, 45

[AHDBV06] Jose Ignacio Alvarez-Hamelin, Luca Dall'Asta, Alain Barrat, and Alessandro Vespignani. k-core decomposition: a tool for the visualization of large scale networks. *Advances in Neural Information Processing Systems 18, Canada*, page 41, 2006. 13, 39, 40

[AHDBV08] Jose Ignacio Alvarez-Hamelin, Luca Dall'Asta, Alain Barrat, and Alessandro Vespignani. k-core decomposition of Internet graphs: hierarchies, self-similarity and measurement biases. *Networks and Heterogeneous Media*, 3:371, 2008. 91, 94

[ALWD05] David Alderson, Lun Li, Walter Willinger, and John C. Doyle. Understanding Internet topology: principles, models, and validation. *IEEE/ACM Transactions on Networking (TON)*, 13(6):1205–1218, 2005. 38

[AVG$^+$99] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra. Routing Policy Specification Language (RPSL). RFC 2622, IETF, June 1999. 31

[BA99] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, (1999). 19, 35, 44, 50

[BBGW04] Michael Baur, Ulrik Brandes, Marco Gaertler, and Dorothea Wagner. Drawing the AS graph in 2.5 dimensions. In János Pach, editor, *Graph Drawing, New York, 2004*, pages 43–48. Springer, 2004. 53

[BEH+07] Giuseppe Di Battista, Thomas Erlebach, Alexander Hall, Maurizio Patrignani, Maurizio Pizzonia, and Thomas Schank. Computing the types of the relationships between autonomous systems. *IEEE/ACM Transactions on Networking (TON)*, 15(2):267–280, 2007. 32

[BHM+07] R. Bush, J. Hiebert, O. Maennel, M. Roughan, and S. Uhlig. Testing the reachability of (new) address space. In *Proceedings of the 2007 SIGCOMM workshop on Internet network management (INM'07)*, 2007. 90

[BJ07] A. Banerjee and J. Jost. Spectral plot properties: Towards a qualitative classification of networks. In *European Conference on Complex Systems*, October 2007. 61

[Bra89] R. Braden. Requirements for Internet Hosts – Communication Layers. RFC 1122, IETF, October 1989. 27

[BT02] T. Bu and D. Towsley. On distinguishing between Internet power law topology generators. In *Proceedings of IEEE Infocom 2002*, New York, NY, June (2002). 44, 47, 50, 51

[But06] Steve Butler. Lecture notes for spectral graph theory. *Lectures in Nankai University, Tianjin, China*, 2006. 71, 85

[But07] Steve Butler. Interlacing for weighted graphs using the normalized laplacian. *Electronic Journal of Linear Algebra*, 16:90–98, 2007. 87

[CCG+02] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power laws in Internet topologies revisited. In *Proceedings of IEEE Infocom 2002*, New York, NY, June 2002. 35, 37

[CD00] J. M. Carlson and John Doyle. Highly optimized tolerance: Robustness and design in complex systems. *Physical Review Letters*, 84(11):2529–2532, 2000. 38

[CDGT88] D.M. Cvetković, M. Doob, I. Gutman, and A. Torgašev. *Recent Results in the Theory of Graph Spectra*. North-Holland, 1988. 54

[CDZ97] Kenneth L. Calvert, Matthew B. Doar, and Ellen W. Zegura. Modeling Internet topology. *IEEE Communications Magazine*, 35(6):160–163, (1997). 42, 70, 71

[CFSV06] V. Colizza, A. Flammini, M.A. Serrano, and A. Vespignani. Detecting rich-club ordening in complex networks. *Nature Physics*, 2(2):110–115, 2006. 53

[CGJ+04] Hyunseok Chang, Ramesh Govindan, Sugih Jamin, Scott J. Shenker, and Walter Willinger. Towards capturing representative AS-level Internet topologies. *Computer Networks*, 44(6):737–755, 2004. 37, 66

[CGW89] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Pseudo-random graphs. *Combinatorica*, 9(4):345–362, 1989. 73

[CHK+06] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. MEDUSA - New Model of Internet Topology Using k-shell Decomposition. *ArXiv Condensed Matter e-prints*, 2006. 39

[Chu97] Fan R. K. Chung. *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics)*. American Mathematical Society, 1997. 54, 70, 73, 86

[CJW03] Hyunseok Chang, Sugih Jamin, and Walter Willinger. Internet connectivity at the AS-level: an optimization-driven modeling approach. In *MoMeTools '03: Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research*, pages 33–46, Karlsruhe, Germany, 2003. 41

[CJW06] Hyunseok Chang, Sugih Jamin, and Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In *Proceedings of IEEE Infocom 2006*, Barcelona, Spain, 2006. 33

[CM05] Aaron Clauset and Cristopher Moore. Accuracy and scaling phenomena in Internet mapping. *Physical Review letters, 94:018701*, 2005. 29

[DAHB+06] Luca Dall'Asta, Ignacio Alvarez-Hamelin, Alain Barrat, Alexei Vazquez, and Alessandro Vespignani. Exploring networks with traceroute-like probes: theory and simulations. *Theoretical Computer Science*, 355(1):6–24, 2006. 30

[DF07] B. Donnet and T. Friedman. Internet topology discovery: a survey. *IEEE Communications Surveys and Tutorials*, 2007. to appear. 26

[DKF+07] Xenofontas Dimitropoulos, Dmitri Krioukov, Marina Fomenkov, Bradley Huffaker, Young Hyun, kc claffy, and George Riley. AS relationships: inference and validation. *SIGCOMM Computer Communications Review*, 37(1):29–40, 2007. 26, 33

[Doa96] Matthew B. Doar. A better model for generating test networks. In *IEEE GLOBECOM'96*, London, UK, November (1996). 43, 50, 70

[DRFC06] B. Donnet, P. Raoult, T. Friedman, and M. Crovella. Deployment of an algorithm for large-scale topology discovery. *IEEE Journal on Selected Areas in Communications (JSAC)*, 24(12):2210–2220, 2006. 29

[DW87] J.E. Dennis and D.J. Woods. Optimization in microcomputers: The nelder-meade simplex algorithm. In A. Wouk, editor, *New Computing Environments: Microcomputers in Large-Scale Computing*, pages 116–122. SIAM, 1987. 75

[Eco05] Nicholas Economides. The economics of the Internet backbone. *NYU, Law and Research Paper No. 04-033; and NET Institute Working Paper No. 04-23*, June 2005. 97

[ELR+96] D. Estrin, T. Li, Y. Rekhter, K. Varadhan, and D. Zappala. Source Demand Routing: Packet Format and Forwarding Specification (Version 1). RFC 1940, IETF, May 1996. 26

[ER85] P. Erdös and A. Rényi. On random graphs. In *Mathematical Institute Hungarian Academy, 196*, London, (1985). 34, 50

[FB05] Nick Feamster and Hari Balakrishnan. Detecting BGP configuration faults with static analysis. In *NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation*, pages 43–56, Berkeley, CA, 2005. USENIX Association. 48

[FFF99] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of ACM SIGCOMM 1999*, pages 251–262, Cambridge, Massachusetts, United States, 1999. 29, 34, 36, 37, 49, 50, 96

[FKP02] Alex Fabrikant, Elias Koutsoupias, and Christos H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the Internet. In *ICALP '02: Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, pages 110–122, London, UK, 2002. Springer-Verlag. 38

[FMM+04] Anja Feldmann, Olaf Maennel, Z. Morley Mao, Arthur Berger, and Bruce Maggs. Locating Internet routing instabilities. In *Proceedings of ACM SIGCOMM 2004*, pages 205–218, Portland, OR, 2004. 51

[GALM08] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening Internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *Proceedings of Passive and Active Measurement Conference (PAM)*, April 2008. 84, 92, 95

[Gao01] Lixin Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking (TON)*, 9(6):733–745, 2001. 31

[GMZ03] C. Gkantsidis, M. Mihail, and E. Zegura. Spectral analysis of Internet topologies. In *Proceedings of IEEE Infocom 2003*, San Francisco, CA, April 2003. 18, 38, 71, 85

[GT00] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for Internet map discovery. In *Proceedings of IEEE Infocom 2000*, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE. 26, 27

[GW02] Timothy G. Griffin and Gordon Wilfong. On the correctness of iBGP configuration. *SIGCOMM Computer Communications Review*, 32(4):17–29, 2002. 31

[HAA+07] Bradley Huffaker, Dan Andersen, Emile Aben, Matthew Luckie, k Claffy, and Colleen Shannon. The Skitter AS Links Dataset, 2001-2007. 88

[Hal97] Bassam Halabi. *Internet Routing Architectures*. Cisco Press, 1997. 47

[Han07] Sean Hanna. Representation and generation of plans using graph spectra. In *6th International Space Syntax Symposium*, Istanbul, Turkey, 2007. 71, 85

[HFJ+08] Hamed Haddadi, Damien Fay, Almerima Jamakovic, Olaf Maennel, Andrew W. Moore, Richard Mortier, Miguel Rio, and Steve Uhlig. Beyond node degree: evaluating AS topology models. Technical Report UCAM-CL-TR-725, University of Cambridge, Computer Laboratory, July 2008. 75

[HFU+08] Hamed Haddadi, Damien Fay, Steve Uhlig, Andrew Moore, Richard Mortier, Almerima Jamakovic, and Miguel Rio. Tuning topology generators using spectral distributions. In *Lecture Notes in Computer Science, Volume 5119, SPEC International Performance Evaluation Workshop*, Darmstadt, Germany, 2008. Springer. 72, 84, 85

[HKYH02] P. Holme, B.J. Kim, C.N. Yoon, and S.K. Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):298–305, 2002. 53

[HPM+02] Bradley Huffaker, Daniel Plummer, David Moore, , and k Claffy. Topology discovery by active probing. In *SAINT-W'02: Proceedings of the 2002 Symposium on Applications and the Internet (SAINT) Workshops*, page 90. IEEE Computer Society, 2002. 27, 47

[HPSS03] Oliver Heckmann, Michael Piringer, Jens Schmitt, and Ralf Steinmetz. On realistic network topologies for simulation. In *MoMeTools '03: Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research*, pages 28–32, New York, NY, USA, (2003). 71

[JB06] Shudong Jin and Azer Bestavros. Small-world characteristics of Internet topologies and implications on multicast scaling. *Computer Networks*, 50(5):648–666, 2006. 40

[JRT04] Sharad Jaiswal, Arnold L. Rosenberg, and Don Towsley. Comparing the Structure of Power-Law Graphs and the Internet AS Graph. In *ICNP '04: Proceedings of the Network Protocols, 12th IEEE International Conference on (ICNP'04)*, pages 294–303. IEEE Computer Society, 2004. 39

[JU07] A Jamakovic and S Uhlig. On the relationship between the algebraic connectivity and graph's robustness to node and link failures. In *3rd EuroNGI Conference on Next Generation Internet Networks*, Trondheim, Norway, (2007). 71, 85

[KF75] A. N. Kolmogorov and S. V. Fomin. *Introductory Real Analysis*. Dover Publication, 1975. 103, 104

[KFC$^+$05] Krishnamurthy, Michalis Faloutsos, Marek Chrobak, Li Lao, Jun-Hong Cui, and Allon G. Percus. Reducing large Internet topologies for faster simulations. In *IEEE/IFIP NETWORKING*, pages 328–341, 2005. 37

[LAWD04] Lun Li, David Alderson, Walter Willinger, and John Doyle. A first-principles approach to understanding the Internet's router-level topology. In *Proceedings of ACM SIGCOMM 2004*, pages 3–14, Portland, OR, 2004. 38

[LBCX03] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proceedings of IEEE Infocom 2003*, San Francisco, CA, April 2003. 29, 90

[LCMF08] Y. Li, J.-H. Cui, D. Maggiorini, and M. Faloutsos. Characterizing and Modelling Clustering Features in AS-Level Internet Topology. In *Proceedings of IEEE Infocom 2008*, April 2008. 85

[LH01] Bin Luo and E.R. Hancock. Structural graph matching using the EM algorithm and singular value decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(10):1120–1136, October 2001. 85

[LM08] M. Latapy and C. Magnien. Complex network measurements: Estimating the relevance of observed properties. In *Proceedings of IEEE Infocom 2008*, April 2008. 84

[Mal93] G. Malkin. Traceroute Using an IP Option. RFC 1393, IETF, January 1993. 25, 51

[Mer95] R. Merris. A survey of graph laplacians. *Linear and Multilinear Algebra*, 39:19–31, 1995. 54

[MFM$^+$06] Wolfgang Muhlbauer, Anja Feldmann, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Building an AS-topology model that captures route diversity. *SIGCOMM Computer Communications Review*, 36(4):195–206, 2006. 33, 101

[MH05] Damien Magoni and Mickael Hoerdt. Internet core topology mapping and analysis. *Computer Communications*, 28(5):494–506, 2005. 13, 27, 28

[MHK$^+$07] Priya Mahadevan, Calvin Hubble, Dmitri Krioukov, Bradley Huffaker, and Amin Vahdat. Orbis: rescaling degree correlations to generate annotated Internet topologies. *SIGCOMM Computer Communications Review*, 37(4):325–336, 2007. 41

[MKF$^+$06] Priya Mahadevan, Dmitri Krioukov, Marina Fomenkov, Xenofontas Dimitropoulos, k c claffy, and Amin Vahdat. The Internet AS-level topology: three data sources and one definitive metric. *SIGCOMM Computer Communication Review*, 36(1):17–26, (2006). 37, 48, 51, 52, 58, 59

[MKFV06] Priya Mahadevan, Dmitri Krioukov, Kevin Fall, and Amin Vahdat. Systematic topology analysis and generation using degree correlations. In *Proceedings of ACM SIGCOMM 2006*, pages 135–146, Pisa, Italy, 2006. 41, 47, 55

[MLMB01] Alberto Medina, Anukool Lakhina, Ibrahim Matta, and John Byers. BRITE: an approach to universal topology generation. In *IEEE MASCOTS*, pages 346–353, Cincinnati, OH, USA, August (2001). 36, 71, 80

[MMB00] Alberto Medina, Ibrahim Matta, and John Byers. On the origin of power laws in Internet topologies. *SIGCOMM Computer Communications Review*, 30(2):18–28, 2000. 35

[MQWZ05] Z. Morley Mao, Lili Qiu, Jia Wang, and Yin Zhang. On AS-level path inference. In *SIGMETRICS '05: Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 339–349, 2005. 33

[MRWK03] Z Morley Mao, Jennifer Rexford, Jia Wang, and Randy H. Katz. Towards an accurate AS-level traceroute tool. In *Proceedings of ACM SIGCOMM 2003*, pages 365–378, Karlsruhe, Germany, 2003. 31, 49, 51

[MUF+07] Wolfgang Muhlbauer, Steve Uhlig, Bingjie Fu, Mickael Meulle, and Olaf Maennel. In search for an appropriate granularity to model routing policy. In *Proceedings of ACM SIGCOMM 2007*, Kyoto, Japan, August 2007. 33, 101

[New02] M.E.J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20):871–898, 2002. 52

[NJW02] A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: analysis and an algorithm. In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*. MIT Press, 2002. 72, 85

[NLCK05] Boaz Nadler, Stephane Lafon, Ronald Coifman, and Ioannis Kevrekidis. Diffusion maps, spectral clustering and eigenfunctions of Fokker-Planck operators. In *Neural Information Processing Systems*, 2005. 72

[NM65] J.A. Nelder and R. Mead. A simplex method for function minimization. *Comput. J.*, 7:308–313, 1965. 75

[OPW+08] Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. In search of the elusive ground truth: The Internet's AS-level connectivity structure. In *ACM SIGMETRICS*, Annapolis, USA, June 2008. 94, 96, 101

[OZZ07] Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. Observing the evolution of Internet AS topology. In *Proceedings of ACM SIGCOMM 2007*, Kyoto, Japan, August 2007. 52, 84, 91, 101

[PG98] Jean-Jacques Pansiot and Dominique Grad. On routes and multicast trees in the Internet. *SIGCOMM Computer Communications Review*, 28(1):41–50, 1998. 29

[Pos81] J. Postel. Internet Control Message Protocol. RFC 0792, IETF, September 1981. 26

[Poy63] Pentti Poyhonen. A tentative model for the volume of trade between countries. *Weltwirtschaftliches Archive*, 90:93–100, 1963. 33

[PSV04] Romualdo Pastor-Satorras and Alessandro Vespignani. *Evolution and Structure of the Internet: A Statistical Physics Approach*. Cambridge University Press, 2004. 40

[Quo05] Bruno Quoitin. Topology generation based on network design heuristics. In *CoNEXT'05: Proceedings of the 2005 ACM conference on Emerging network experiment and technology*, pages 278–279, Toulouse, France, 2005. 45

[RLH06a] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271, IETF, January 2006. 25

[RLH06b] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271, IETF, January 2006. 47

[RTM08] Matthew Roughan, Simon Jonathan Tuke, and Olaf Maennel. Bigfoot, sasquatch, the yeti and other missing links: what we don't know about the as graph. In *IMC '08: Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pages 325–330, New York, NY, USA, 2008. ACM. 66

[Sab66] G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966. 53

[SARK02] Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, and Randy H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proceedings of IEEE Infocom 2002*, June 2002. 32, 84

[SBCC98] J. Stewart, T. Bates, R. Chandra, and E. Chen. Using a Dedicated AS for Sites Homed to a Single Provider. RFC 2270, IETF, January 1998. 47

[SFFF03] Georgos Siganos, Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. Power laws and the AS-level Internet topology. *IEEE/ACM Transactions on Networking (TON)*, 11(4):514–524, 2003. 35

[SLH06] L. Shyu, S-Y. Lau, and P. Huang. On the search of Internet AS-level topology invariants. In *Proceedings of IEEE Global Telecommunications Conference (GLOBECOM) 2006*, pages 1–5, San Francisco, CA, USA, 2006. 84

[SMA03] Neil Spring, Ratul Mahajan, and Thomas Anderson. The causes of path inflation. In *Proceedings of ACM SIGCOMM 2003*, pages 113–124, Karlsruhe, Germany, 2003. 42

[SMW02] Neil Spring, Ratul Mahajan, and David Wetherall. Measuring ISP topologies with rocketfuel. In *Proceedings of ACM SIGCOMM 2002*, pages 133–145, 2002. 25, 26, 27

[SR03] A.J. Seary and W.D. Richards. Spectral methods for analyzing and visualizing networks: an introduction. In *Dynamic Social Network Modeling and Analysis*, pages 209–228. National Academic Press, 2003. 72

[SS05] Yuval Shavitt and Eran Shir. DIMES: let the Internet measure itself. *SIGCOMM Computer Communications Review*, 35(5):71–74, 2005. 18, 28

[SW04] G. Sagie and A. Wool. A clustering approach for exploring the Internet structure. In *23rd IEEE Convention of Electrical and Electronics Engineers in Israel*, pages 149–152, 2004. 39

[Syl52] J J Sylvester. A demonstration of the theorem that every homogeneous quadratic polynomial is reducible by real orthogonal substitutions to the form of a sum of positive and negative squares. In *Philosophical Magazine IV*, pages 138–142, 1852. 74

[TGJ+02] Hongsuda Tangmunarunkit, Ramesh Govindan, Sugih Jamin, Scott Shenker, and Walter Willinger. Network topology generators: degree-based vs. structural. In *Proceedings of ACM SIGCOMM 2002*, pages 147–159, Pittsburgh, PA, (2002). 48, 50, 54, 70

[Tho87] A. G. Thomason. Pseudo-random graphs. *Random Graphs '85, North-Holland Mathematical Study*, 144:307–331, 1987. 73

[TMSV03] Renata Teixeira, Keith Marzullo, Stefan Savage, and Geoffrey M. Voelker. In search of path diversity in ISP networks. In *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pages 313–318, Miami Beach, FL, 2003. 29

[TSGR04] Renata Teixeira, Aman Shaikh, Tim Griffin, and Jennifer Rexford. Dynamics of hot-potato routing in ip networks. *SIGMETRICS Perform. Eval. Rev.*, 32(1):307–319, 2004. 101

[VHE02] Danica Vukadinovic, Polly Huang, and Thomas Erlebach. On the spectrum and structure of Internet topology graphs. In *IICS '02: Proceedings of the Second International Workshop on Innovative Internet Computing Systems*, 2002. 71, 85

[Wax88] Bernard M. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications (JSAC)*, 6(9):1617–1622, December (1988). 43, 49, 50, 70

[WJ02] Jared Winick and Sugih Jamin. Inet-3.0: Internet topology generator. Technical report, University of Michigan Technical Report CSE-TR-456-02, 2002. 51, 71, 85

[WL06] X. Wang and D. Loguinov. Wealth-based evolution model for the Internet AS-Level topology. In *Proceedings of IEEE Infocom 2006*, April 2006. 73, 84

[Woo97] D. R. Wood. An algorithm for finding a maximum clique in a graph. *Operations Research Letters*, 21(7):211–217, January 1997. 54

[WS04] A. Wool and G. Sagie. A clustering approach for exploring the Internet structure. *Proceedings of 23rd IEEE Convention of Electrical and Electronics Engineers in Israel*, pages 149–152, Sept. 2004. 85

[YJB02] Soon-Hyung Yook, Hawoong Jeong, and Albert-Laszlo Barabasi. Modeling the Internet's large-scale topology. *Applied Physical Sciences*, 99(21):13382–13386, October 2002. 39

[ZCD97] Ellen W. Zegura, Kenneth L. Calvert, and Michael J. Donahoo. A quantitative comparison of graph-based models for Internet topology. *IEEE/ACM Transactions on Networking (TON)*, 5(6):770–783, (1997). 49, 70

[Zho06] Shi Zhou. Characterising and modelling the Internet topology, the rich-club phenomenon and the PFP model. *BT Technology Journal*, 24, 2006. 44, 47, 50, 51, 81, 84

[ZW05] P. Zhu and C. Wilson. A study of graph spectra for comparing graphs. In *The 16th British Machine Vision Conference*, September 2005. 103

[ZZZ07] S. Zhou, G.-Q. Zhang, and G.-Q. Zhang. Chinese Internet AS-level topology. *IET Communications*, 1(2):209–214, April 2007. 51