Modeling the Topology of the Internet

An Assessment

Pascal Habegger and Hanspeter Bieri

Institute of Computer Science and Applied Mathematics
Neubrueckstrasse 10, 3012 Bern, Switzerland

July 2002

Abstract

It is well known that the Internet has experienced a fascinating growth in the past. Unfortunately, precise measurements and statistics of its size and relevant descriptions of its structure are not available. It is the goal of ongoing research to get a better understanding of the Internet topology at different levels of abstraction (e. g. at the router-level or AS-level). This would help to design more realistic simulation scenarios when evaluating the performance of network protocols. It would alleviate the search for structural weaknesses in the Internet, would help the design of next-generation Internet protocols, and might even permit to predict the further evolution of the Internet in the near future. This Technical Report gives an assessment of current research results by means of listing and evaluating different important models which attempt to explain the formation of the topology of the Internet as we experience it today. It also contains a survey of major topology generators which are currently used by the network simulation community. It concludes with an outlook at future work and open problems.

Keywords: **Internet topology, topology generator, Internet models**

Contents

1 Introduction

Since its commercialization in 1995[1], the growth of the Internet has been tremendous. According to NetSizer[2], the Internet roughly consisted of an average of 153,810,000 hosts in January 2002 compared to 69,970,000 hosts in January 2000—the size of the Internet had more than doubled within two years. As the Internet is not centrally administrated or operated, no one is capable of giving exact numbers of its current size or its topological structure. Fundamental questions like *What does the Internet look like?* or *How has the network structure of the Internet evolved in the past few years?* are therefore difficult to answer.

There exists no comprehensive description of the Internet, which has grown more or less uncontrolled in the past few years. In particular the question concerning the structure of the Internet topology cannot be answered. Although snapshots of the Internet topology have been taken in regular time intervals for the past four years, this data is incomplete and probably gives a wrong picture. The lack of complete information can be largely attributed to the fact that no single administrative center exists which monitors and controls the world-wide growth of the Internet. The Internet grows peripherally, due to a multiplicity of external influences which are only partly known. Thus, insufficient information on Internet topologies makes it difficult to understand the complex dynamics driving the growth of the Internet.

It is not only scientific curiosity which is pushing the investigation of the Internet to provide answers to the aforementioned questions. Due to the large scale of the Internet and to the missing central management, many interesting problems relating to the Internet have emerged. For instance, *routing* in the Internet has become more and more complicated. Routing instabilities caused by temporary router failure or overload, incorrect router configuration, or loss of link connectivity are getting more and more frequent, resulting in parts of the Internet being temporarily cut off. The Internet has also become more *vulnerable to failures and attacks* leading to large parts of the Internet being disconnected. For instance, when the Internet cable SEA-ME-WE-3 connecting Southeast Asia, the Middle East and Western Europe, was damaged in November 2000, large parts of the Internet ceased to function. The breakdown of the cable, which stretches about 39,000 km and which is one of the world's largest and fastest undersea telecommunications cables, was possibly caused by a ship's anchor truncating the cable. It took several days to repair, leaving most hosts located on the Australian continent unreachable from the remaining parts of the Internet. In order to locate and adjust these structural deficiencies of the Internet, a better understanding of the Internet topology and the rules of its growth is required.

Another field of research which would profit from a deeper understanding of the Internet topology is the *design* and *evaluation* of Internet applications and communication protocols. As it could be shown that distributed algorithms (e. g. routing algorithms) are often highly sensitive to the underlying network topology, care has to be taken when designing new algorithms for the Internet. Moreover, modified or new applications need

---

[1] In 1995, the NSFnet was dismantled by the NSF (National Science Foundation) and replaced by a commercial Internet backbone ([Zak02]). At the same time, the NSF implemented a new backbone, called very high-speed Backbone Network Service (vBNS), and reverted back to a research network. Originally, the NSFnet replaced ARPANET as the main government network linking universities and research facilities.

[2] http://www.netsizer.com, Telcordia Technologies.

to be tested using network simulation prior to deploying them in the Internet ([FP01]). With the advance of computer hardware and the availability of powerful simulation tools (e. g. the *network simulator* ns-2 [ns02], OPNET[3], or the *Scalable Simulation Framework* SSF[4]), simulation of large networks has become manageable. However, in order to obtain significant simulation results, "realistic" network topologies have to be used in simulation projects. "Realistic" means that these networks should display topological properties similar to those found in the Internet requiring a better knowledge of the Internet topology.

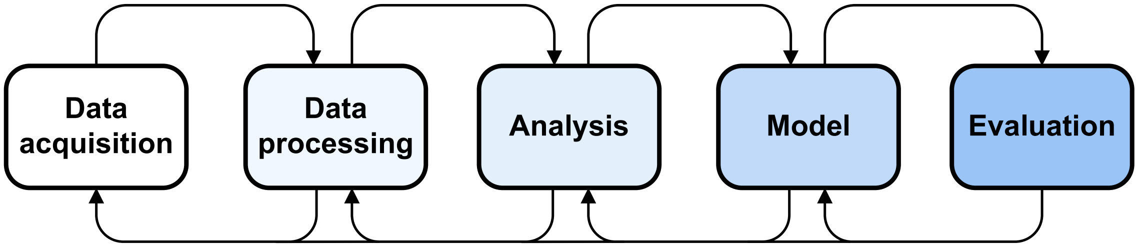

Figure 1: Iterative modeling process

The *iterative* modeling approach taken by most researchers is illustrated in Figure 1. After the acquisition and the processing of snapshot data of the Internet topology, the result is analyzed for significant topological properties. Based on the results, a model of the Internet is then proposed which requires testing and evaluation in a next step. As no complete snapshots of the Internet are available and new topological properties are being discovered over time, existing models need to be refined (or possibly rejected) and re-evaluated.

There exist several kinds of topology data which can be obtained from public providers like ISPs and universities. Sometimes, this data needs to be preprocessed in some way or possibly be improved by merging the data available from different sources. Beside the problem of acquiring topology data, the statistical analysis of the data is of great importance. Usually, it is done by converting the topology data to an undirected graph which can then be analyzed using the numerous mathematical instruments employed in graph theory. Finally, the proposed model needs to be evaluated by comparing it to a series of snapshot data of the Internet taken at different points in time.

Although the study of the Internet topology has attracted more and more attention within the research community during the past few years, no model exists to date which succeeds in explaining the reasons why the Internet has evolved as it has. Many different approaches have been proposed in the past decade aimed at modeling and/or understanding the Internet. There are *flat random*, *hierarchical*, and *degree-based* models, which attempt to produce topologies similar to the Internet. Also, different *evolutionary models* have been introduced which are meant to imitate the evolution of the Internet topology. Unfortunately, none of these models has succeeded so far, and it is questionable whether there will ever exist such a model. It is not known what forces and influences are driving the growth of the Internet. A model attempting to closely imitate the growth of the Internet would have to identify and quantify at least the most significant ones. All

---

[3] `http://www.opnet.com`.

[4] `http://www.ssfnet.org`.

known models fail to do that. For this reason, a new class of models is needed which better imitates the complex growth processes of the Internet.

It is the goal of our report to analyze and compare the major approaches which have been proposed by the research community during the past few years. Special emphasis will be given to the individual steps which make up the iterative modeling process as illustrated in Figure 1. Interestingly, not all proposed models have their roots in the Internet research. Models for dynamic systems also exist in other areas of research, e. g. in mathematics, sociology, economics, physics, biology, or chemistry. Recently, some of these approaches have been applied to Internet topologies. Thus Internet topology research has become of interest to a larger, interdisciplinary group of scientists from very different fields of research. So-called topology generators are used to produce network topologies which exhibit similar properties as found in the Internet. Our report lists and assesses the major generators for each topology model.

The report is organized as follows: Extensive background information regarding the Internet and its topology is given in Section 2. In Section 3 the modeling process, which includes data retrieval and processing, analysis, design, and evaluation of a new Internet model, is presented in more detail together with related terms taken from graph theory. Section 4 discusses various reasons motivating the investigation of the Internet topology. Since no complete information on the topology of the Internet is available, other ways have to be found in order to obtain representative structure information of the network. This is looked at in Section 5. In Section 6 we report important statistics of the Internet topology of the past and present which will then be used to evaluate prevalent topology models listed in Section 7. For each model, a list of common topology generators is given followed by a brief assessment of the respective modeling approach. Our report concludes with final remarks in Section 8.

2 The Internet

What exactly is the Internet? What is it composed of? Questions like these are not easily answered and require a certain degree of background information which will be given in this section. For lack of space, our report does not deal with the history of the Internet. A brief overview of the evolution of the Internet is given in *Hobbes' Internet Timeline* [Zak02]—a more detailed account can be found in [Nau00].

From a macroscopical point of view, the Internet is composed of thousands of so-called autonomous systems (AS) of varying size and importance which are themselves computer networks. Each of these computer networks consists of (possibly thousands of) hosts, e. g. computers, routers, or switches, and links connecting these hosts. Together they make up the very *fabric* of the Internet. The ASes on their part are interconnected according to service agreements set up between selected pairs of ASes. These agreements govern the way data traffic is routed between neighboring ASes.

Partly due to the absence of a central administration of the Internet, no complete topological map of the Internet is available. For this reason, numerous mapping projects have been conducted in the past in order to provide a better picture of the structure of the Internet. The goals of these projects are diverse: some studies try to obtain a better spatial map of the space occupied by the devices and other physical objects of the

Internet (e. g. hosts and routers), others target at mapping and exploring the Cyberspace[5] induced by the Internet. For instance, the connectivity of the Internet, the World Wide Web (WWW), and virtual worlds which have emerged on top of the Internet (e. g. MUDs[6] or peer-to-peer networks) have been the object of various extensive mapping projects ([GT00, HFMC01, DK01]).

2.1 Autonomous systems

The structure of the global Internet corresponds to a loose compound of half-autonomous networks, or domains, which offer a specific set of services with their own pricing policy to their respective customers. Domains vary in number of hosts, geographic size, and function. So-called *provider domains* offer their transmission and switching facilities as a service for data exchange to other domains, e. g. a campus or a corporate network. These provider domains usually interconnect at so-called neutral Internet exchange points and can vary in the geographical scope of their operations from regional, to national and international.

For instance, at CIXP (CERN Internet eXchange Point) in the Geneva area, Switzerland, the following national and international networks, ISPs[7] and research institutes interconnect (selection only):

- **Research institutes and networks:** CERN (European Organization for Nuclear Research), GÉANT (formerly TEN-155), Institut national de physique nucléaire et de physique des particules (IN2P3), SWITCH (Swiss Academic and Research network)
- **Swiss ISPs:** Sunrise/diAx, CABLECOM, OnSpirix, Swisscom IP-Plus Internet Services, VTX
- **International ISPs and Telecom providers:** AT&T, Cablecom, Cable & Wireless (Petrel Communications), COLT, Deutsche Telekom, Easynet, France Telecom, Global Crossing, GTS (Ebone), infonet, KPNQwest (EUnet), Worldcom (MCI, UUNET)

Each of these domains represents or is part of a so-called *autonomous system* (AS). RFC[8] 1930 defines an AS as a “connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy.”[9] Each AS maintains traffic routing contracts with its neighboring ASes which governs the way data traffic is routed between them. Four kinds of relationships between ASes are common ([Hus99a, Hus99b, Gao00]):

Customer-provider: The customer AS pays the provider AS for access to the rest of the Internet. The provider offers a transit service to its customers.

Peering: Two ASes agree to exchange traffic between their respective customers free of charge.

---

[5] Dodge and Kitchin define *Cyberspace* as “the conceptual space within Information and Communication Technologies” ([DK01]).

[6] Multi-User Dungeons.

[7] Internet Services Providers.

[8] Request for Comments.

[9] Sometimes, an AS is defined as a set of routers under a single technical administration. This definition is ambiguous since it does not mention what interior gateway protocols are used.

Mutual-transit agreement: Two ASes provide access to the rest of the Internet for each other.

Mutual-backup agreement: Two ASes provide backup connectivity to the Internet for each other in case the primary provider of one of the ASes fails.

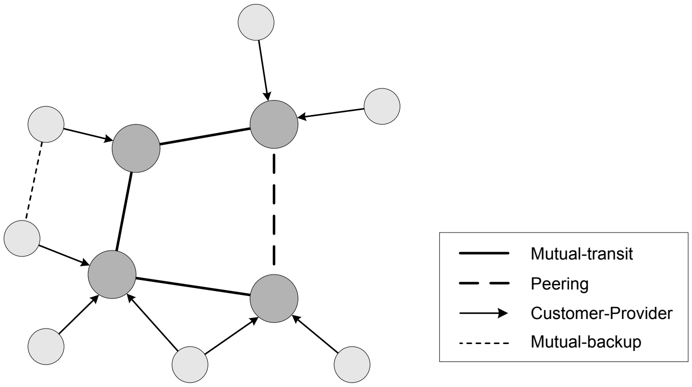

Figure 2: Relationships between 11 ASes

In Figure 2 a sample network is given which illustrates the possible relationships between neighboring ASes.

In the interior of an AS, the administration of the AS realizes its own routing policy using an *Interior Gateway Protocol* (IGP). Typically, the protocols RIP[10] or OSPF[11] are used. The main purpose of the *Exterior Gateway Protocols* (EGP) consists of the exchange of routing information between neighboring domains (or ASes) of the Internet (so-called inter-network routing). Today's de-facto standard in the Internet is the Border Gateway Protocol (BGP), currently in its forth version (cf. [RL95]). The ISO Inter-Domain Routing Protocol (IDRP, ISO 10747) is another scalable inter-autonomous system routing protocol which will support IPv4 as well as the next generation of IP (IPv6). It is expected that IDRP will be adopted in the Internet when BGP becomes obsolete. A brief overview of BGP will be given in the next subsection. Figure 3 illustrates the use and the relationship of the individual protocols for inter-domain and intra-domain routing.

Each AS has a 16 bit Autonomous System Number (ASN) assigned to it (i. e. 0–65,535) which is globally unique.[12] Currently three Regional Internet Registries (RIR) in different territories of the earth are responsible for these assignments:

- RIPE[13] coordinates administrative and technical tasks within Europe and North Africa.
- ARIN[14] manages the administration of North and South America, the Caribbean and Subsaharan Africa.

---

[10] Routing Information Protocol.

[11] Open Shortest Path First Protocol.

[12] Currently, efforts are made in the IETF (Internet Engineering Task Force) working group for Inter-Domain Routing (idr) to make provisions for a 4-byte ASN (`http://www.ietf.org/html.charters/idr-charter.html`).

[13] Réseaux IP Européens, `http://www.ripe.net`.

[14] American Registry for Internet Numbers, `http://www.arin.net`.

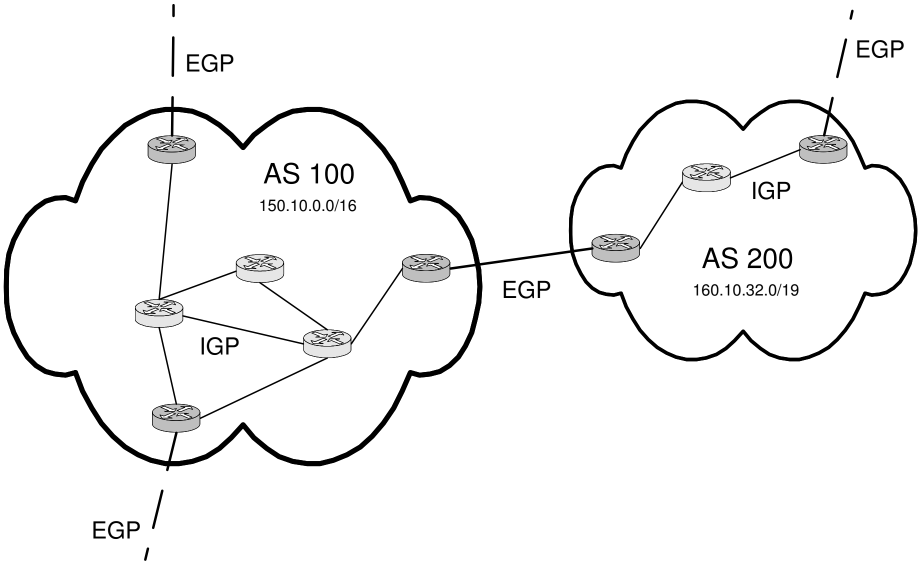

Figure 3: Interior/Exterior Gateway Protocols

- APNIC[15] represents the Asia Pacific region.

An ISP or an international company may have several ASNs making the task of finding all ASNs belonging to a certain company rather tedious and sometimes even impossible. Although many ISPs own many ASNs no central database exists which helps to identify such relationships. The current situation is additionally aggravated by continuous mergers among ISPs.

2.2 Border-gateway protocol

As mentioned in the previous section, BGP is a protocol for the exchange of routing information between routers in a network of ASes like the Internet. Each BGP-enabled router keeps a routing table containing all networks which can currently be reached by that router. For each destination network a path of ASNs (so-called ASPATH) is stored which is used to reach the destination. Figure 4 shows parts of a sample routing table[16]. The meaning of the individual columns of this table is as follows:

Network holds the address of the destination network in CIDR notation (see below). Since most networks can be reached using different paths, several entries on separate lines for the same destination network can be found in a routing table.

Next Hop holds the IP address of the next system which will be used to forward data packets to the destination network.

Metric, LocPrf, Weight can be used to rate a path to the destination AS. Thus, the router can select the most effective path for routing in terms of cost.

---

[15] Asia Pacific Network Information Centre, `http://www.apnic.net`.

[16] This output was produced on the router with the command `show ip bgp`.

```
   Network          Next Hop            Metric LocPrf Weight Path
*  192.138.32.0/19  128.177.255.6                          0 3557 6461 568 i
*                   128.177.255.5                          0 3557 6461 568 i
*                   206.220.240.223                        0 10764 1 1239 568 i
*                   203.181.248.233                        0 7660 1 1239 568 i
*                   167.142.3.6                            0 5056 3561 568 i
*                   193.140.0.1                            0 8517 9000 2548 568 i
*                   193.0.0.56                             0 3333 1103 6453 1239 568 i
*                   134.55.20.229                          0 293 568 i
*                   192.121.154.25                         0 1755 1800 1239 568 i
*                   204.29.239.1                           0 6066 3549 1239 568 i
*                   195.219.96.239                         0 8297 6453 1239 568 i
*                   216.140.14.127          91             0 6395 568 i
*                   216.140.8.63            54             0 6395 568 i
*>                  134.24.127.30           57             0 1740 568 i
```

Figure 4: Part of a BGP routing table (from Oregon Exchange BGP route viewer)

Path contains an ordered list of ASes which have to be traversed in sequence in order to reach the destination AS. This column is of particular interest for the following studies of BGP routing tables since they describe the links between ASes.

The latest and most widespread version of BGP, BGP-4, implements *Classless Inter-Domain Routing* (CIDR, [FLYV93]) which offers the possibility to easily expand the IP address space as used by the currently used Internet protocol (IPv4).

According to IPv4 each IP address corresponds to a 32 bit number identifying both the sending and the receiving host. An IP address is made up of two parts:

1. **Network identifier**
   Identifies a specific network of the Internet
2. **Host identifier**
   Identifies a specific host in a network

As the networks heavily vary in their size (i. e. the number of contained hosts) the addressing scheme of the Internet protocol provides four classes (or addressing templates):

- **Class A**
  For large networks with many hosts (7 bit network address, 24 bit host address)
- **Class B**
  For middle-sized networks (14 bit network address, 16 bit host address)
- **Class C**
  For small networks with less than 256 hosts (21 bit network address, 8 bit host address)
- **Class D**
  Reserved for multicasting (28 bit multicast address)

The inflexible partitioning of the address space led to an accelerated shortage of available IP numbers in the past ([Hus01]). For instance, a medium-sized company comprising little more than 254 hosts was forced to buy a class B address which left most of the allocated address space (one class B address can hold up to 65,533 IP addresses) unused. For this reason, the overall available address space grew swiftly scarce. With the introduction of CIDR the problem of wasting large chunks of address space could be solved, albeit not permanently.

Basically, CIDR eliminates the concept of class A, B, and C networks and replaces it with a generalized “network prefix” (often also termed “IP prefix”). CIDR allows the delegation of pieces of what used to be called “network numbers” to customers, and therefore makes it possible to utilize the available address space more efficiently. As an example, the CIDR address `192.138.32.0/19` consists of the IP address (`192.138.32.0`) and a mask length (`19`). The mask length specifies the number of leftmost contiguous significant bits in the corresponding IP address. In this example, the network `192.138.32.0` with IP-prefix `19` contains 32 class C networks.

CIDR can be used to perform route aggregation where a single route can cover the address space of several “old-style” network numbers and thus replace a lot of old routes. This lessens the local administrative burden of updating external routing, saves routing table space in all backbone routers and reduces route flapping (rapid changes in routes), and thus CPU load in all backbone routers.

This approach was first proposed in 1992 as a scheme called *supernetting* ([FLYV93]). The concept of a supernet can be compared to the public telephone system (POTS, Plain Ordinary Telephone Service) where many phone numbers are aggregated to a group using area codes. This kind of aggregation corresponds to the supernetting of many networks of the Internet using CIDR—currently network prefixes of 12 bits up to 24 bits are in use.

To be better able to coordinate the configuration of individual BGP routers in the Internet and the routing between them, the Internet Routing Registry (IRR, `http://www.irr.net`) emerged early in 1995, at a time when providers worldwide were preparing for the end of the NSFnet Backbone Service and the birth of the commercial Internet. Designed to be a next-generation database development effort with participants from many international networking organizations, its data, which is available free of charge, may be used by anyone worldwide to help debug, configure, and engineer Internet routing and addressing. Originally, the IRR comprised only five databases but has grown since to a large collection including more than 50 different databases. Its databases contain many existing routing policies between ASes of the Internet. Each policy is specified using the Routing Policy Specification Language (RPSL, [AVG+99]) which enables a network operator to specify routing policies at various levels in the Internet hierarchy, in particular at the AS level.

Although routing data and policies contained in these registries are publicly accessible, no guarantees can be made as to whether the stored information is up-to-date, complete, or accurate. It is assumed that the routing implementations in the networks reflect the announced policies, i. e. the router configurations in a network are consistent with the published policies. No policing of the registered routing policies is performed. Thus, routing policies stored in these registries may differ from observed routing policies in the Internet. For these reasons, analysis and research of the Internet topology (especially at the AS level) have been mainly based on real routing data measured at selected backbone BGP routers in the past. Many of those BGP routing table dumps are available free of charge in the Internet, and are usually generated over a longer period of time. Further information on selected sources of routing data will be given in Section 5.

2.3 The topology of the Internet

It has to be emphasized that there exists more than one topology[17] of the Internet. The Internet looks very different at the router level than at the AS level. At the router level it is made up of millions of hosts, routers, switches, and other networking devices which are joined by physical connections or wireless links. When looking at the Internet at the AS level however, it consists of thousands of autonomous systems whose interconnections are defined by bilateral commercial agreements expressed in service contracts. Also, the Internet has been rapidly growing in the past decade. Not only has the size of its topology increased but also its structure has significantly changed. Therefore, there exist different topologies for different points in time.

In our report, only the Internet topology at the AS level will be of interest. This topology is often represented as an *unweighted*, *undirected*, and *simple* graph $G = (V, E)$ where the set V stands for its vertices or nodes (i. e. the individual ASes), and the set E represents its edges (i. e. network connections between two ASes). For the rest of this report, this graph will be referred to as the *AS connectivity graph*, or *AS graph.* Some authors also use a directed, possibly labelled graph when describing the Internet at the AS level. When discussing prominent modeling approaches in Section 7, different definitions of the AS graph used by the authors will be explicitly noted.

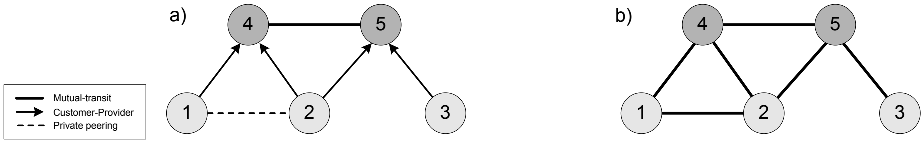

Figure 5: AS connectivity does not imply reachability

Undirected, unweighted graphs may only be used when studying the *interconnection* structure of the Internet. As the routing between ASes is done on the basis of the BGP which is a policy-based routing protocol, connectivity does not imply reachability ([Gao00]). This is illustrated as an example in Figure 5. Figure 5 a) shows aa network of five ASes with different bilateral service agreements. For the customer AS 3 to reach AS 1, data traffic has to go through ASes $3, 5, 4, 1$. Only ASes 4 and 5 have a mutual-transit agreement; the remaining ASes do not offer any transit or public-peering agreements. Reachability information is lost when looking at the corresponding connectivity graph shown in Figure 5 b). According to this graph, AS 1 can be reached from AS 3 using several different routes, for instance when going through ASes $3, 5, 2, 1$. Obviously, this is an illegal routing path since the autonomous systems 1 and 2 do not offer public peering.

Most modeling approaches analyzed in this report focus on AS connectivity graphs. This is mostly due to the fact that AS reachability information is not as easily available as AS connectivity data. However, several heuristics have been proposed lately which deduce

---

[17] Webopedia "defines" the term *topology* as follows: *"The shape of a local-area network (LAN) or other communications system. Topologies are either physical or logical."* (`www.webopedia.com`)

AS service agreements from ASPATH information available in BGP (cf. [Gao00, CCG+02]). In Section 7.2 these heuristic methods will be presented in more detail.

2.4 Selected terms from graph theory

Since the topology of the Internet is usually described as a graph, this subsection gives selected terms from graph theory. The following terminology, largely taken from [GY99, Wat99], will be used throughout the rest of this report. In our report, every graph is undirected unless noted differently.

loop
A loop (sometimes also denoted as *self-loop* [GY99]) is an edge that joins an endpoint to itself.

simple
A graph is *simple* if it has neither loops nor multi-edges.

degree
The *degree* (or *valence*) $\mathrm{deg}(v)$ of a vertex v is the number of edges incident on v plus twice the number of loops.

neighborhood
The *neighborhood* $\Gamma(v)$ of a vertex v is the subset that consists of the vertices adjacent to v (not including v itself).

path
In a graph, a path from vertex v_0 to vertex v_n is an alternating sequence

$$W =< v_0, e_1, v_1, e_2, \ldots, v_{n-1}, e_n, v_n >$$

of vertices and edges, where the endpoints of the edge e_i are the vertices v_{i-1} and v_i.
The *length* of a path is the number of edges in the path.

distance
The *distance* between two vertices $u, v \in V$, $d(u, v)$ is the length of the shortest path between u and v.

connected
A graph is connected if for every pair of vertices u and v, there is a path from u to v.

diameter
The diameter of graph G is given by

$$\mathrm{diam}(G) = \max_{u,v \in V} \{d(u, v)\}$$

clique
A subset S of V of a graph $G = (V, E)$ is called a clique if every pair of vertices in S is joined by an edge, and no proper superset of S has this property.

eigenvalue and **spectrum**

The eigenvalues of a graph are defined as the eigenvalues of its adjacency matrix. λ is called an eigenvalue of a square matrix A if there exists a non-zero vector $\vec{v}$ (called an eigenvector of A) such that

$$A\vec{v} = \lambda\vec{v}$$

The set of eigenvalues of a graph is called the spectrum of the graph.

3 The modeling process

Analyzing and understanding the topology of the Internet is an iterative multistage process, which is illustrated in Figure 6 (compare also with Figure 1):

1. In a first step, representative snapshots of Internet topology data are retrieved. As no complete data is available, heuristic methods have to be used for data acquisition (see Section 5).

2. After that, the snapshot data needs to be processed in order to build a representation of the Internet topology, usually in the form of an *AS connectivity graph*. Some authors have also proposed heuristic methods for retrieving AS reachability information from snapshot data. A weighted or labelled graph can be constructed, sometimes also referred to as the *annotated AS graph* ([Gao00]).

3. Using statistical analysis, the resulting AS graph has to be examined for meaningful graph invariants and dynamical growth properties (see Section 6). These results lay the basis for the definition of a new topology model.

4. Based on the computed graph invariants and related topological attributes of the snapshot data, a parameterized model of the Internet can be defined. This represents the most critical step of the modeling process. Afterwards, topologies produced with the newly developed model have to be compared to those of the "real" Internet (see Section 7).

The remaining parts of our report are dedicated to this iterative multistage process and will describe the individual stages in more detail. To begin with, the next section presents the different reasons motivating the analysis and the modeling of the Internet topology.

4 Applications for Internet models

Being able to identify major topological properties of the Internet helps to gain a better understanding of the complex dependencies among the individual hosts and networks of the largest global network of networks.

The following goals are the driving force behind the Internet topology research (see also Figure 6):

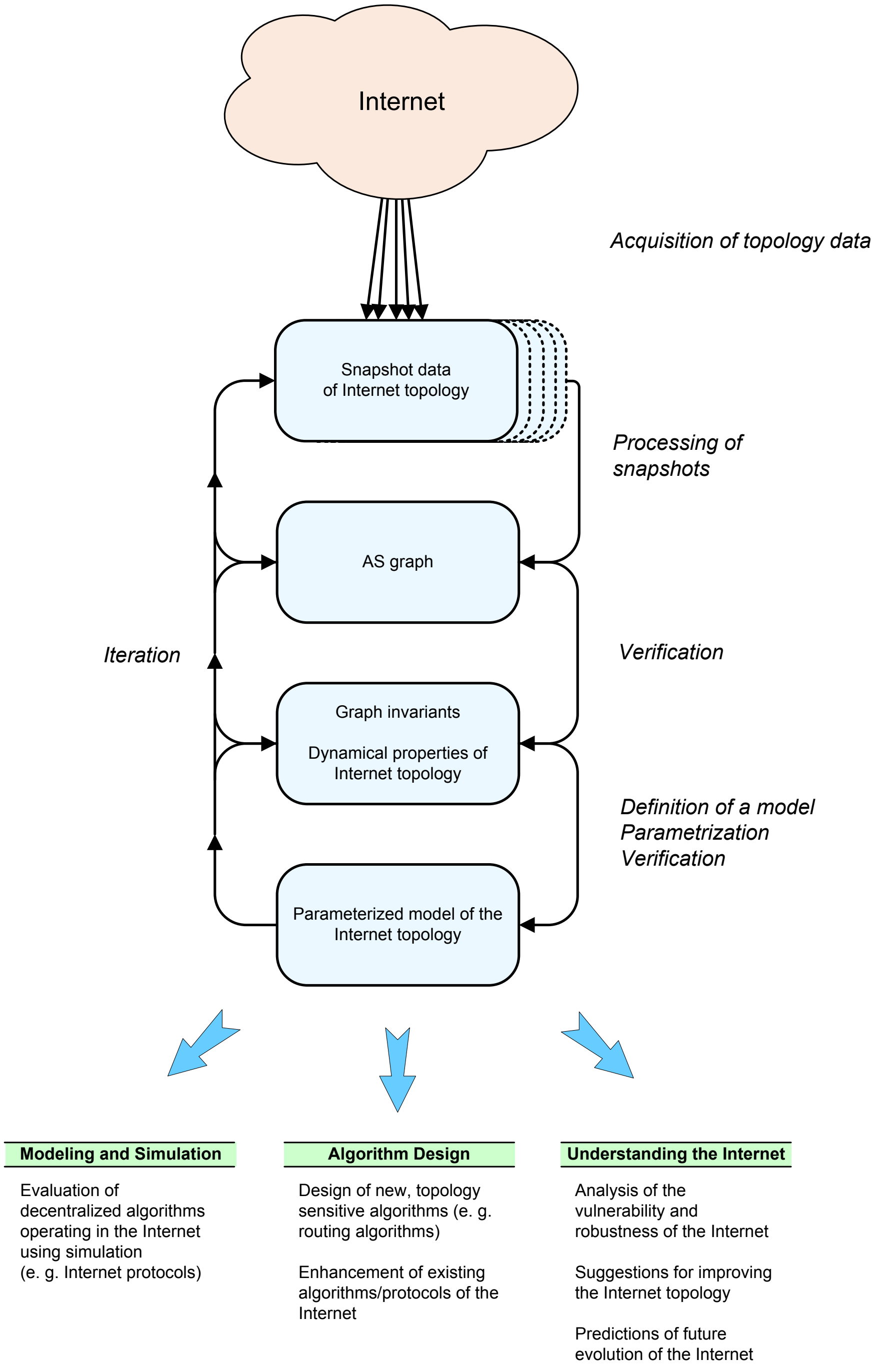

Figure 6: Modeling process and applications

- Since new network protocols are thoroughly tested in simulation prior to deploying them in the Internet, the use of more realistic topologies results in more meaningful simulations and evaluations of the protocols. It can be shown that network protocols perform in a significantly different way for different network topologies.
- If the underlying topology is known, more efficient decentralized algorithms can be designed which take advantage of the unique topological properties of the Internet. Particularly, routing algorithms, search strategies, or applications based on multicasting profit from it.
- A thorough understanding of the evolution of the Internet topology and its status quo allows for making qualitative statements about its robustness and vulnerability. In a limited way, it also permits to predict the future evolution of the Internet and may offer suggestions for improvements of the topology.

The following subsections elaborate on these goals and provide several published examples.

4.1 Modeling and simulation

Before using modified or newly developed protocols or distributed algorithms in the Internet on a large scale, reliable predictions have to be obtained regarding their realtime performance and scalability. Simulation represents a powerful means for making those predictions, and is used and accepted in the research community today ([BEF$^+$00, FP01]). As recent research has repeatedly revealed ([RTY$^+$00, MP01b]), the choice of the underlying network topology is predominant when realizing a simulation project. Depending on the topological properties of the network structure, the gained simulation results may heavily vary and be different from measurements made in the Internet after the deployment of the respective protocol. Although the topology has no effect on the correctness of a network protocol, it can have a major impact on its performance.

Many of today's Internet protocols are highly sensitive to the network topology they are deployed in. When designing new communication protocols for the Internet, special attention has to be given to scalability and resilience issues. Protocols tested for small homogenous networks with well-known topology and reliable connection links perform completely different from protocols for very large heterogenous networks with incomplete connectivity information and unreliable links. Several studies have been conducted investigating the impact of network structure on the design and the performance of multicast protocols:

- Radoslavov et al. carried out four case studies dealing with selected multicast protocols ([RTY$^+$00]) and the impact of topology on protocol performance. For example, they studied a receiver-driven alternate path routing (APR) protocol to route around congested links in a multicast tree. They define two metrics *success rate* (probability of finding an uncongested path to the source) and *extra path length* (additional length of the uncongested path) and show in simulation that these metrics are significantly affected by the underlying network topology.
- Doar et al. evaluated naïve multicast routing in several simulation runs and with different random and hierarchical network topologies ([DL93]). They observed a

significant performance degradation in random graphs compared to more realistic network topologies.

- Magoni et al. ([MP01b]) compare their agent search algorithm, which is very sensitive to the underlying network topology, to the expanding rings search (ERS) algorithm which is implemented in existing protocols such as YAM and QoSMIC. They define four variables of interest: bandwidth usage, the number of optimal agents found, the number of attempts made to find an optimal agent, and efficiency. Using simulation and different network topologies generated by various topology generators (see Section 7), they found that the kind of topology model used for protocol simulation has a crucial impact on the simulation results. They also stated that even the latest topology generators do not yet produce realistic topologies.

The topology of networks used in simulation projects need to display similar, realistic properties as found in the real Internet. For performance reasons the size of generated networks in simulations is usually several orders of magnitude smaller than the Internet, although efforts exist to deploy distributed and parallel simulation frameworks (e. g. [CNO99]) when using large models of the Internet (> 15,000 nodes).

Since the understanding of the Internet topology has increased over time, numerous so-called *topology generators* have been proposed in the last few years with the aim to produce graphs with the discovered or assumed properties. At first, popular topology generators produced simple random graphs [Bol01, Wax88] which were then used in simulation projects. More recently, more “realistic” generators have been developed based on different approaches (e. g. [Doa96, ZCB96, ZCD97, MMB00, JCJ00]). In Section 7 a brief review of selected generators will be presented.

4.2 Algorithm design

An analysis of the inter-domain topology and route stability can have an impact on the design of future *routing protocols*, the distribution of routing information in the Internet, as well as the implementation and provisioning of routers ([GR97]). Route instability is typically caused by temporary router failure or overload, router misconfiguration, or loss of link connectivity ([LMJ99]). Instable routes can increase route computation overhead on routers. A higher degree of connectivity between domains often leads to an increase of alternate paths which may necessitate advanced routing algorithms in the Internet.

Furthermore, a better knowledge of the underlying topology of the Internet helps to improve existing algorithms or open new ways of designing more efficient strategies for achieving certain tasks in the Internet or similar networks (e. g. searching in adhoc networks [ALPH01], or the prevention of distributed Denial-of-Service attacks [PL00]).

4.3 Understanding the Internet

Many complex systems display a surprisingly high degree of tolerance against random failures. Although this property is often attributed to the redundant wiring of the Internet, Albert et al. point out in [AJB00] that it is not shared by all redundant systems. It is only displayed by a class of inhomogeneously wired networks, so-called *scale-free* networks. However, error tolerance comes at a high price as these networks are extremely vulnerable to attacks. If an informed agent eliminates the most connected nodes, the highest damage in terms of failure can be inflicted to the Internet. In this way, the structure of the Internet, guaranteeing a high degree of redundancy, becomes also its Achilles heel when under attack. As test results have shown ([AJB00]), taking out 2.5 % of the nodes with the highest connectivity may lead to a triplication of its diameter—often, a fragmentation into several non-connected sub-networks may occur. This behavior is not shared when randomly removing 2.5 % of all nodes. The authors attribute this behavior to the scale-free characteristic of the Internet and its evolutionary development.

Knowing the major defects of today's Internet topology, future enhancements of the network can reduce the risk of (partial) failure of the network ([GR97]). In-depth study of the global Internet may also show possible ways of protecting the Internet against willful attacks.

Undenyingly, the ability to predict the further evolution of the Internet in the coming two, three, or even five years would be of great value to most developers of Internet software. Topology sensitive algorithms like routing protocols, multicasting, (audio and/or video) streaming, IP telephony, or large-scale peer-to-peer software could profit from a more precise knowledge of the future topology of the Internet.

In contrast to this, unforeseen social or technological breakthroughs can change the shape of the topological Internet in a dramatic way within a few years, making it inherently impossible to reliably predict its future evolution. A new "killer application" may come along taking the place of existing WWW applications and changing the shape of current Internet traffic. For instance, several peer-to-peer networks like Gnutella, Morpheus, or eDonkey have appeared lately which may become even more important in the future. Other examples are real-time applications such as telephony, video, or (massive) multiplayer gaming [FP01]. For these reasons, available long-term conjectures of the future size of the Internet or of other selected properties should be met with a certain amount of scepticism. However, predictions for the near future (e. g. one to two years) are more likely to turn out to be true.

As an example, Faloutsos et al. predicted in 1999 for the beginning of the year 2000 that the size of the Internet graph at the AS-level would amount to 6,364 nodes and 13,576 edges assuming a 45% increase of nodes ([FFF99]). On January 2 2000, according to [Nat02, ANTC02], the Internet incorporated 6,474 nodes and 13,895 edges on that day, which approximately matched the prediction of the year before. The conjecture for the year 2002 was a little bit more off: 13,380 nodes (NLANR: 12,518) and 29,421 edges (NLANR: 28,765). Since Faloutsos et al. based their work on data available from [Nat02], the same source was used by us for verifying their predictions. However, it has to be taken into account that the topology data available from [Nat02] is by no means complete—it only contains a large part of the real AS graph. Nevertheless, the three power-laws describing the Internet topology, which were discovered by Faloutsos et al. in 1999, seem to hold, at least for the near future. A more in-depth view of these power-laws will be given in Section 7.4.

The following sections discuss the individual aspects of the modeling process illustrated in Figure 6. In a first step, several ways of acquiring topology data of the Internet are described. This data will then be processed resulting in an AS graph.

5 Acquiring and processing snapshots of the Internet topology

It is extremely difficult to accurately analyze the topology of the Internet. There exists no complete map of the inter-domain topology since the Internet is geographically distributed and is not managed by one global administration. Also, it is not possible to obtain complete BGP routing tables from all existing ASes as most ASes are not willing to publicly share their routing tables for security and commercial reasons.

Therefore, other ways have to be chosen in order to construct a representative picture of the topology of the Internet. However, these snapshots are usually incomplete as they represent only a (albeit large) part of the Internet. The following section lists publicly available sources of information on current or past topology data of ASes. It then shows in what ways the retrieved data is processed in order to obtain an AS connectivity graph.

5.1 Routing tables of BGP routers

The most commonly chosen way of acquiring topology data of the Internet is by parsing full routing table dumps of BGP routers ([FFF99, JCJ00, MBB00, MP01a, CCG+02]). These large routing tables (usually between 40 MB and 350 MB of text files) represent a partial snapshot of the global AS map periodically taken during a certain time span. They capture (a large part of) the topological structure of the Internet at the AS level but fail to reflect the routing policies deployed by the individual ASes.

The University of Oregon Route Views Project has made BGP routing tables publicly available for research purposes since 1997. The Route Views router, `route-views.oregon--ix.net`, uses multi-hop BGP peering sessions with backbones at interesting locations. It uses AS6447 in its peering sessions, and routes received from neighbors are never passed on nor used to forward traffic. `route-views.oregon-ix.net` itself does not announce any prefixes. Routing table data is available from `archive.routeviews.org` back to April 2001. For the period of November 1997 till mid March 2001 BGP routing table dumps can be retrieved online from the National Laboratory for Applied Network Research (NLANR, [Nat02]).

5.2 Update messages between BGP routers

Instead of using instantaneous dumps of BGP routing tables at a router, chronological traces of routing transitions over a longer period of time can be analyzed ([GR97]). An UPDATE message is the primary message used to communicate information between two BGP routers ([Ste99]). When a BGP speaker A advertises a prefix to a neighboring BGP router B or withdraws a previously advertised prefix, an UPDATE message is sent from A to B. For instance, if A is sending an UPDATE message containing a withdrawal of a network prefix to B, this typically means that the path between the router A and that prefix is temporarily down and the prefix cannot be reached.

Govindan et al. evaluated 21-day segments of their traces, parsing the ASPATH field in UPDATE messages ([GR97]). This technique usually provides richer views on inter-domain topologies as more than one ASPATH to an address prefix are seen over the period of the trace which results in denser AS graphs. However, these traces are more difficult to obtain and do not yield complete maps of the Internet topology, similar to BGP router dumps.

5.3 Bottom-up Approach

Using a router-level Internet map, an AS *overlay* can be computed on top of it. Using a BGP routing table as a look-up reference, each router of the map gets an ASN assigned to it. In this way, an AS graph can be derived which represents real routing paths in the Internet according to the initial Internet map. However, this approximate model contains several inaccuracies, as mentioned in [TGSE01]:

1. Not all IP addresses have matching entries in the BGP routing table. No mapping between ASN and those IP addresses can therefore be made. Alternatively, the IRR routing policy database could be used when matching specific IP addresses to ASNs.
2. The collapsing algorithm used for computing the AS overlay produces many disjoint clusters of nodes belonging to the same ASes due to the lack of a complete Internet map and ASN assignment. A heuristics has to be applied to resolve these inconsistencies.

5.4 Looking glass information

Many ASes have so-called looking glasses on their web sites which allow read-only access to their routers (see Figure 7 for an example). In this way, an observer can query the path along which traffic is routed from that AS to another AS. Also, it is possible to retrieve BGP summaries of their peerings with neighboring domains. These tables do not include BGP links with other autonomous systems that are more than one hop away; particularly, no full dumps of BGP routing tables are publicly accessible since this would put too much load on the routers. Therefore, routing information obtained from looking glass sites alone cannot be used to construct significant AS graphs. Nevertheless, these BGP summaries contain complete information on the local connectivity of selected nodes or ASes of the graph which can be used to improve the quality of the produced AS graph (e. g. [CCG+02]).

5.5 Internet Routing Registries

Routing information available from Internet Routing Registries (IRR, cf. Section 2.1) can be used to further augment AS connectivity maps. Although publication of an ISP's routing policy in an IRR is voluntary and has been incomplete in the past ([Gao00]), more and more Internet exchange points force their members to register their routes and peering policies in publicly available routing policy databases ([CCG+02]). Therefore, AS connectivity information obtained from these registries can be used to construct an AS graph. After building an AS graph from BGP routing tables and looking glass data, Chen et al. ([CCG+02]) use RPSL databases to enhance the AS graph.

5.6 Routing Information Service

At the RIPE Network Coordination Centre, one of the three RIRs worldwide mentioned in Section 2.1, a so-called *Routing Information Service* (RIS) is currently being developed.[18]

---

[18] `http://www.ripe.net/ris/`.

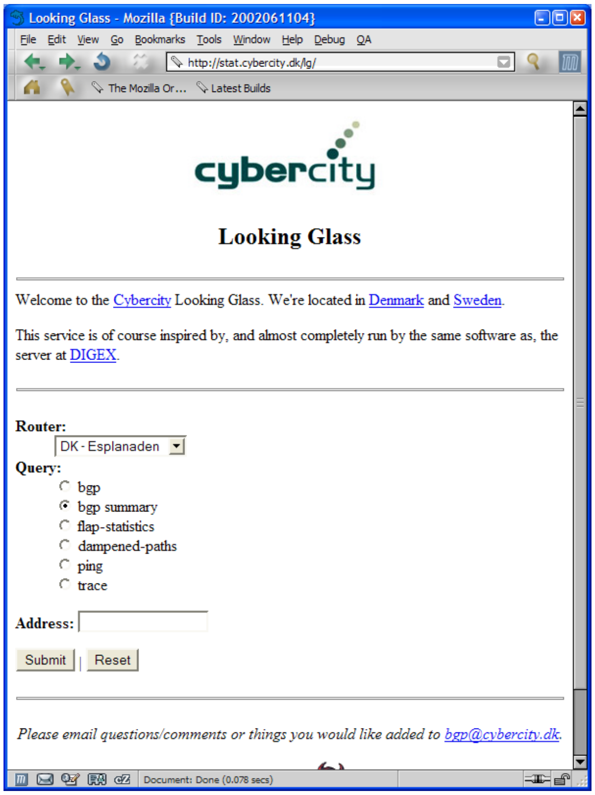

Looking Glass

Welcome to the Cybercity Looking Glass. We're located in Denmark and Sweden.

This service is of course inspired by, and almost completely run by the same software as, the server at DIGEX.

Router: DK - Esplanaden

Query:
- ○ bgp
- ● bgp summary
- ○ flap-statistics
- ○ dampened-paths
- ○ ping
- ○ trace

Address:

Please email questions/comments or things you would like added to bgp@cybercity.dk.

Figure 7: Example of a looking glass website

In the RIS project, a database has been created which stores masses of BGP routing information collected at different locations around the world, particularly at Internet exchange points ([AU99]). At each of those points a so-called *Remote Route Collector* (RRC) is installed which is basically a machine collecting routing BGP information and which is periodically transferring the data to RIPE NCC for further processing. The RRCs are not routing servers, i. e. they do not forward or announce BGP routing information to peers. At RIPE NCC, the data received from the individual RRCs (nine at the time of writing) is integrated into a comprehensive view which can be used to query historical information on the Internet routing from a global point of view.

By retrieving and merging the different BGP routing table dumps available from RIS, a more complete picture of the current Internet topology can be constructed. However, to the authors' knowledge no analysis pertaining to Internet topology research has been done so far based on this data.

All approaches for acquiring Internet snapshot data mentioned in this section only supply incomplete information on the AS connectivity graph. By merging connectivity data acquired from different sources the resulting AS graph can be further augmented ([CCG+02]). However, the question remains unsettled whether these snapshots convey a representative picture of the real AS connectivity graph.

The processing of the acquired snapshot data of the Internet results in an AS graph as is illustrated in Figure 3. In the subsequent modeling step, the graph needs to be investigated for significant topological properties.

6 Statistical analysis

In this section major quantitative properties of current and past AS connectivity graphs are presented as a means of reference for the ensuing evaluation of the existing modeling approaches of Internet topologies (cf. Section 7).

Topology information was extracted from BGP routing tables made available by the Route Views Project as described in 5.1. The AS graphs were constructed from BGP routing tables using ANT[19] and analyzed using GINT[20]. Both tools, written in Perl and in Java respectively, have been developed by the authors and simplify the retrieving and analysis of BGP routing table information. Data was retrieved and analyzed on a monthly basis over a period of four years, starting on November 15, 1997 till November 15, 2001.[21] This timeseries of AS graph data will be referred to as the *RVP data set*.

Several important properties can be computed for undirected graphs. The following list summarizes the most interesting statistical properties of AS graphs measured during the observation period:

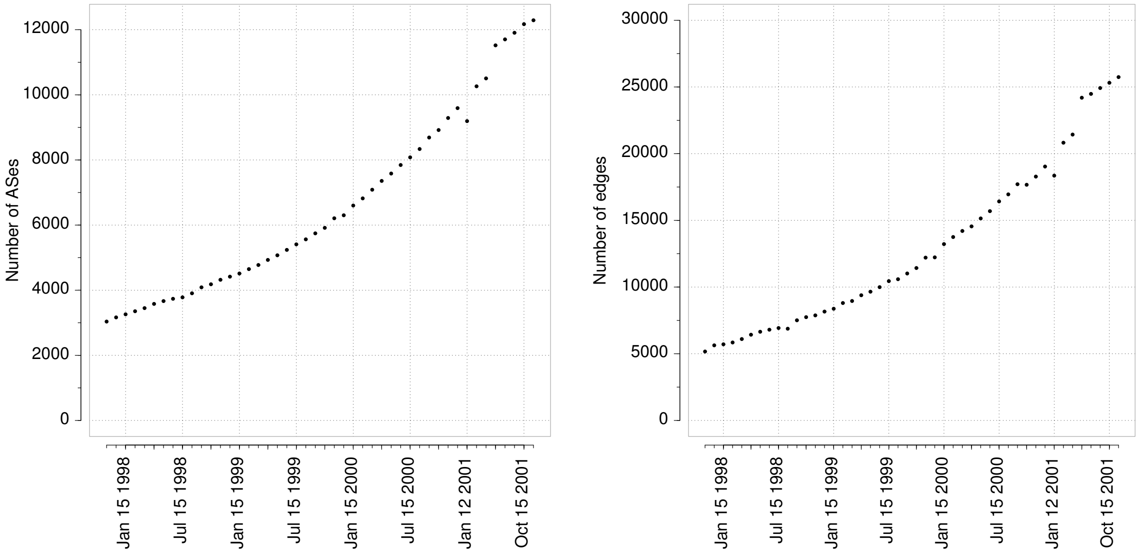

(a) Number of ASes (b) Number of edges

Figure 8: Size of the Internet at the AS level

- Considering the evolution of the size of the Internet at the AS level, Figures 8(a) and 8(b) show that both the number of nodes and edges have more than quadrupled in the past four years. At the same time, the maximum degree of a node has increased from 591 to 2,529 which also represents a quadruplication (Figure 9).
- The average degree of a node has slightly increased from 3.40 to 4.19 which is an indication that the Internet has become better interconnected.

---

[19] Analyzer of Network Topologies.

[20] Generator of INternet Topologies.

[21] Unfortunately, no routing tables were available at the Route Views Project for the period of April 2001 to June 2001.

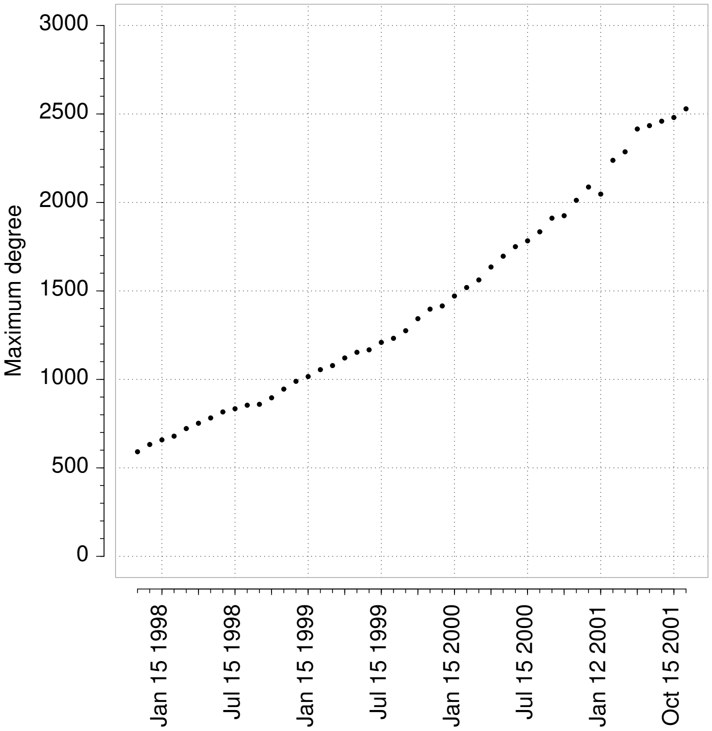

Figure 9: Maximum degree

- The degree distribution of Internet graphs is highly askew: On November 15 2001, according to the BGP routing table dump available from the Route Views Project, 90% of all nodes were of degree 5 or less, 33% of all nodes were of degree 1. The heavy-tailed degree distribution will be of further interest when studying several power-laws of the Internet in Section 7.4.

- Another intriguing fact can be discovered when examining the diameter of the Internet (Figure 10(a)): although size and degree have more than quadrupled during the observation period, our analysis indicates that the diameter of the AS connectivity graph has remained stable, oscillating between 9 and 11. This is particularly noteworthy as there exists no central administration of the Internet coordinating its growth and is an indication that the Internet graph belongs to the category of small-world graphs which will be presented in the next section.

- The *characteristic path length L(G)* is another important graph invariant which is defined by Watts in [Wat99] as follows:

  > The *characteristic path length* $L(G)$ of a graph G is the median of the means of the shortest path lengths connecting each vertex $v \in V$ to all other vertices.

  $L(G)$ represents the typical distance among the distances $d(v_i, v_j)$ between every pair of two distinct vertices v_i and v_j. $d(v_i, v_j)$ denotes here the length of the shortest path between the two vertices v_i and v_j. In Figure 10(b) we have plotted

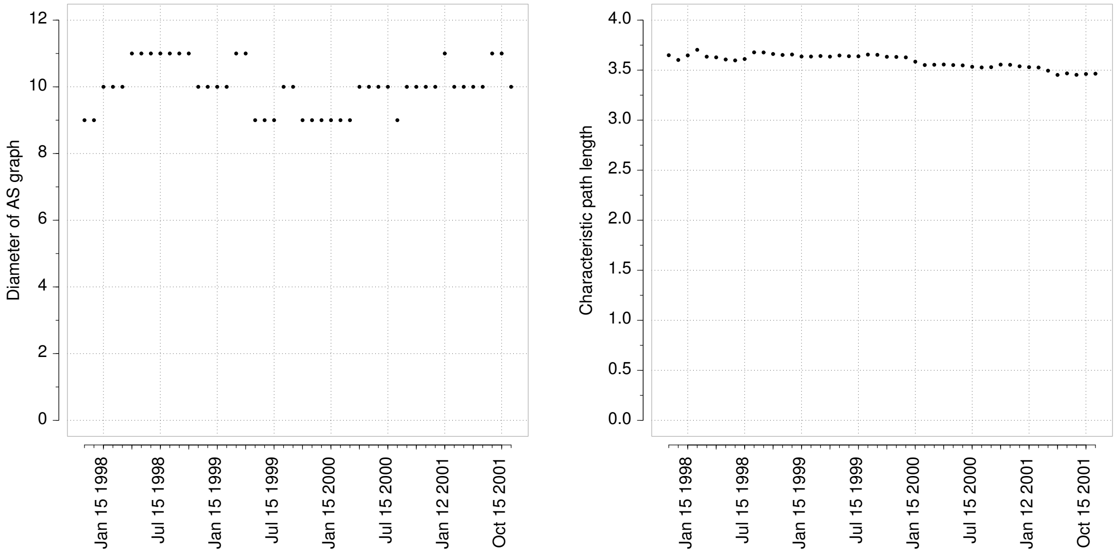

(a) Diameter (b) Characteristic path length

Figure 10: Diameter and path length of the Internet at the AS level

L for the observation period. Again, it can be seen that L has remained relatively small despite the exponential increase of the size of the AS graph.

- A measure for expressing the *clustering* of a graph is given by the so-called *clustering coefficient* γ_v of a vertex v for its neighborhood $\Gamma(v)$. γ_v characterizes the extent to which vertices adjacent to any vertex v are adjacent to each other and is defined as follows ([Wat99]):

  $$\gamma_v = \frac{|E(\Gamma(v)|}{\binom{|V(\Gamma(v)|}{2}}$$

  where $|V(\Gamma(v)|$ and $|E(\Gamma(v)|$ are the number of vertices and edges, respectively, in the neighborhood of v.[22]

  To get a measure of clustering over the entire graph, the clustering coefficient γ of G corresponds to the averaged γ_v computed for all vertices of G.

  All Internet graphs have a relatively high clustering coefficient γ which indicates that the Internet is heavily clustered. Figure 11 shows the development of γ for the past four years. In contrast, random graphs generated according to the model $\mathcal{G}(n, M)$ (see Section 7.1) of similar size have a clustering coefficient below 0.1.

- Each AS graph contains a relatively large maximum clique of size 8 or more[23] (see Figure 12). In contrast, in random graphs of similar size as in the *RVP* data set, only small maximum cliques of sizes 3 to 6 could be found.

---

[22] The definition of γ_v given in [Wat99] is incomplete since γ_v is undefined for vertices $v \in V$ with $\deg(v) \leq 1$. For this reason, we set $\gamma_v := 1$ for all v with $\deg(v) \leq 1$ so that γ_v is defined for all $v \in V$.

[23] Finding the maximum clique of a graph is an NP-complete problem. GINT implements *Algorithm 457* proposed by Bron and Kerbosch for finding all cliques in an undirected graph ([BK73]).

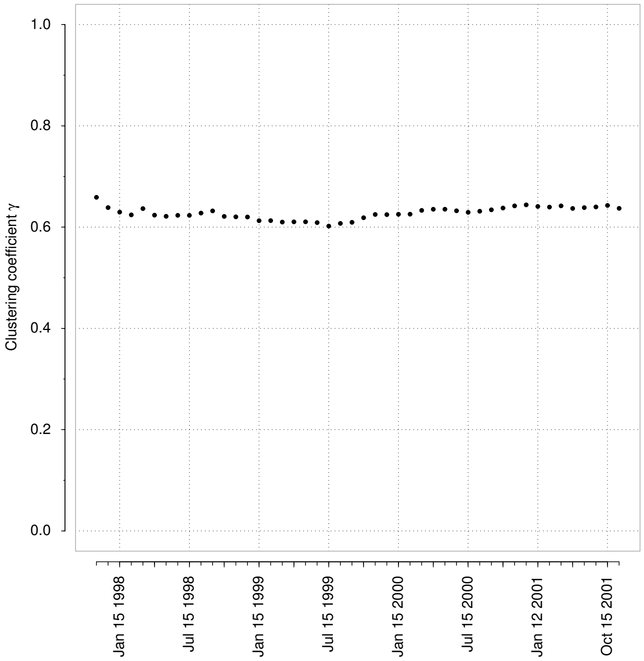

Figure 11: Clustering coefficients of *RVP* dataset

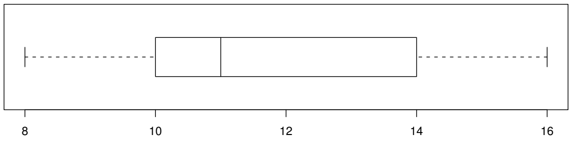

Figure 12: Boxplot of the maximum clique size

Our statistical analysis shows that the clustering coefficient, the diameter, and the characteristic path length of the AS connectivity graphs have remained stable during the whole observation period. In addition, we have found relatively large cliques in the AS graph and observed an overall increase of clique size over time. We are convinced that these facts should be considered when designing an Internet topology model.

Numerous Internet models have been proposed in the literature which rely on some of the statistical results presented here. In the following section, major models are presented which attempt to explain and/or reproduce the topological structure of the Internet.

7 Modeling approaches

During the past decade, numerous models have been proposed with respect to Internet topology research. These approaches can be classified in two different categories:

- *Descriptive models* are based on major graph invariants, like average node degree,

degree distribution, diameter, and so forth. They often contain algorithms which synthetically generate graphs exhibiting those graph properties. However, they do not offer any or only little explanation why Internet graphs display these topological properties. Most flat random methods and power-law graphs belong to this category.

A prominent subgroup is represented by *degree-based models*. They contain algorithms for the generation of Internet topologies which are parametrized based on the degree distribution of a graph. Several current topology generators like INET or PLRG take the degree distribution into account when generating topologies.

- *Explanatory models* attempt to explain the structure of the Internet and its growth. Invariants of the Internet graph are then used for assessing their approach, although this is often done in a limited way. Hierarchical models and the approaches of Barabási and Albert are major representatives of this category.

Topology generators are used to synthetically produce topologies displaying desired structural properties. Table 1 gives an overview of available Internet models and the topology generators which implement the respective models.

	BRITE	GT-ITM	INET	NEM	PLOD	PLRG	TIERS
Flat random [cf. 7.1]	○	○		●			○
Small-world [cf. 7.3]							
Powerlaw-based [cf. 7.4]			●	●	●	●	
Spectrum [cf. 7.6]							
Hierarchical [cf. 7.2]	○	●					●
Evolutionary models [cf. 7.5]	○			○			

Table 1: Popular topology generators and models they support. Models which are fully (partially) supported by a topology generator are represented by • (○).

The following sections provide a chronologically ordered overview of major approaches, starting with the oldest ones. For each modeling approach, known topology generators are listed which are (or have been) used by network researchers in simulation projects or for studying the growth of the Internet. Some topology generators implement more than one model—in this case, attention will only be given to models not yet implemented by previous topology generators.

7.1 Flat random methods

At least four different types of methods using random graphs can be found in the literature. All these approaches have in common that a graph $G = (V, E)$ is built by adding edges to a given set of nodes V subject to a probability function $P(u, v)$ where u, v are arbitrary nodes of G. No special consideration is given to the structure of the produced graph (e. g. hierarchies and the like). In [ZCD97] this family of approaches is termed *flat random methods*.

7.1.1 Random graphs

In graph theory (cf. [Bol01]), the two most frequently occurring definitions of random graphs are given by the sets $\mathcal{G}(n,M)$ and $\mathcal{G}\{n,P(\text{edge}=p)\}$. The first definition considers all graphs $G=(V,E)$ with the labelled vertex set $V=\{1,2,...,n\}$ and having $0 \le M \le \binom{n}{2}$ edges. Each graph in $\mathcal{G}(n,M)$ has the same probability. $\mathcal{G}\{n,P(\text{edge}=p)\}$ with $0 \le p \le 1$ contains all graphs with vertex set $V=\{1,2,...,n\}$ and in which the edges are chosen independently and with probability p.

Using the *pure random method* ([ZCD97]), topology generators add edges with a fixed probability p which corresponds to the model $\mathcal{G}\{n,P(\text{edge}=p)\}$. This simple approach fails to capture significant properties of Internet-like topologies, as in-depth statistical analysis of produced graphs has yielded ([ZCD97, RTY+00]).

The following sections deal with three other common ways of generating random topologies.

7.1.2 Waxman model

B. M. Waxman proposed in [Wax88] a simple random graph model for the purpose of running several network simulation experiments. In his approach, n nodes are randomly distributed over a rectangular coordinate grid. Then m edges are added between pairs of nodes u, v with the probability $P(u,v)$ that depends on the distance $d(u,v)$ between them:

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

where L is the maximum Euclidean distance between any two nodes, and $\alpha,\beta \in (0,1]$. The parameter β can be used to control the density of the generated graph, larger values of β producing denser graphs. α influences the probability of connecting nodes which are locally close—small values increase the number of short edges.

7.1.3 Other models

E. W. Zegura et al. define in [ZCD97] two other generation methods for creating random graphs. The *exponential* method is a slightly modified version of Waxman's model:

$$P(u,v)=\alpha e^{\frac{-d(u,v)}{L-d(u,v)}}$$

It ensures that the probability of an edge between two nodes u and v approaches zero as the distance between these two vertices approaches L.

The *Locality* method, also to be found in [ZCD97], is basically a natural refinement of the pure random method:

$$P(u,v)=\begin{cases}\alpha & , \quad d(u,v)<r\\ \beta & , \quad d(u,v)\ge r\end{cases}$$

Depending on the Euclidian distance between two vertices and on a constant r, different edge probabilities apply. The Locality method is a special case of the model $\mathcal{G}\{n,p_{ij}\}$, where $0 \le p_{ij} \le 1$ for $1 \le i < j \le n$.

7.1.4 Topology generators

There exist several topology generators which support flat random methods, particularly the model introduced by Waxman. The generators BRITE, GT-ITM, and NEM, which not only support flat random methods, are mentioned in subsections 7.5.3.1, 7.2.3.2, and 7.7 respectively.

ntg by S. Hotz ([HN92]) creates an n-level hierarchy of random graphs. In a first step, all nodes are created and arranged on a plane. Afterwards, pairs of nodes are connected following certain criteria. Source code and $n$tg are not publicly available.

rtg (Random Topology Generator) written by L. Wei and L. Breslau ([WE94]) implements the construction algorithm proposed by Waxman.

7.1.5 Assessment of flat random methods

Although flat random methods are a simple and easy way to generate network topologies these approaches fail to realistically model Internet topologies for different reasons:

- The methods themselves do not guarantee that the generated graphs will be connected which is usually a crucial prerequisite when running network simulations. Further tests and modifications (usually by adding further edges) have to be performed to ensure connectedness of the produced graphs.

  One easy way of restoring the connectedness of a graph consists of finding all components[24] and then connecting vertices of different components.

- The produced graphs do not display all properties of today's Internet. For instance, Waxman's approach does not consider local clustering of nodes or low diameters of large graphs.

- Primarily meant to create topologies at the router-level of the Internet, the methods generate *spatial* graphs relying on coordinate attributes assigned to their nodes. In contrast topologies at the AS-level are *relational* graphs—there is no trivial mapping between AS nodes and coordinates.

Flat random methods generally do not produce large graphs exhibiting major topological properties of the Internet—Waxman used his algorithm only for generating small graphs containing about 25 nodes.

7.2 Hierarchical models

Several models have been proposed in the last decade which explain the structure of the Internet at the AS level using hierarchical models. In the following sections, the most common ones are briefly presented and assessed.

7.2.1 Transit-Stub

In order to provide a more adequate router network model of the Internet than the Waxman model, M.B. Doar ([Doa96]) and E.W. Zegura et al. ([ZCB96]) devised a new graph

---

[24] This can be accomplished for instance by performing a depth-first traversal of the graph.

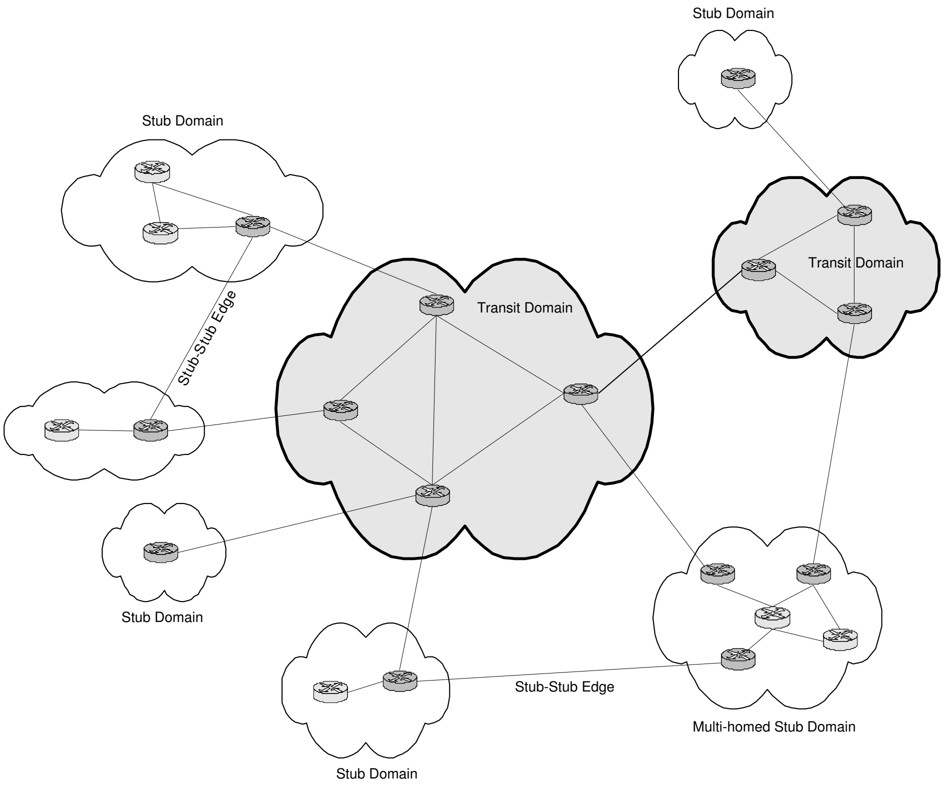

Figure 13: Model of the Internet domain structure

generation method using a more hierarchical domain structure. The authors classified each AS domain in the Internet as either a *stub* domain or a *transit* domain ([ZCD97]). In a stub domain, data traffic between any two nodes u and v goes through that domain only if either u or v is in that domain. In a transit domain, this restriction is dropped. Its main purpose lies in the interconnection of stub domains. An example model of the Internet domain structure is depicted in Figure 13.

Three types of ASes are distinguished:

Transit domain: It comprises a set of backbone nodes which are fairly well connected. Each backbone node also connects to other transit domains or stub domains.

Single-homed stub domain: It is connected via a gateway node to one transit domain. It may also have links to other stubs.

Multi-homed stub domain: It is connected via gateway nodes to several different transit domains and may also contain links to other stubs.

This distinction of domains imposes a hierarchy on the nodes: data traffic originating in the stub domain s_i will only pass through transit domains on its way to the destination stub domain s_j. If $s_i = s_j$, data traffic will not leave that stub domain.

The topology generator GT-ITM, written by E. W. Zegura et al., uses the hierarchy model mentioned above to produce network topologies at the router level of the Internet. Its generation algorithm will be described in subsection 7.2.3.2.

Topologies generated in this way seem to be more characteristic for the Internet than flat random methods ([ZCD97]). Yet, this method fails to produce realistic topologies at the AS level, as important properties of the Internet are lacking (e. g. power-laws do not hold [MMB00, RTY$^+$00]).

7.2.2 Heuristic classification of ASes

Govindan et al. [GR97] postulate a four-level hierarchy of the domains of the Internet classified by function which is depicted in Table 2. They defined four degree ranges based on their studies of several route traces of the Internet between mid 1994 and late 1995. Within a degree range, the fraction of nodes with a given degree remains stable during the observed period of time. Their results also show that the connectivity between domains is significantly non-hierarchical.

Class	Degree range	Fraction of domains	Types of domain
C_1	≥ 28	0.9 %	National or international backbones
C_2	10–27	3.1 %	Large North American regional providers and European national networks
C_3	4–9	9 %	Smaller regional providers, and large metropolitan area providers
C_4	1–3	87 %	Smaller metropolitan area providers and multi-campus corporate or academic networks

Table 2: Degree classification of Internet domains by Govindan et al. [GR97]

Figure 14 shows the degree-based classification of nodes in Internet graphs during the observation period. As can be seen, the fractions of the individual classes defined by Govindan have not changed, although the partitioning of service providers in four classes given in Table 2 does not hold anymore. From the results obtained in [SARK02] it can be seen that a degree-based classification of ASes does not yield a practical distinction of the nodes.

L. Gao presents several heuristic algorithms for inferring AS relationships based on BGP routing table information ([Gao00])). The author proposes an annotated AS graph which is a partially directed graph whose nodes represent ASes and whose edges are classified into four categories: *provider-to-customer*, *customer-to-provider*, *peer-to-peer*, and *sibling-to-sibling*. The author verified inferred AS relationships using real data obtained from AT&T and from WHOIS Lookup services.

L. Subramanian et al. propose in [SARK02] a hierarchical classification algorithm based on a directed AS graph constructed from BGP routing tables available from multiple BGP routers. Their algorithm for classifying AS relationships consists of a pruning process repeatedly applied to the AS graph. The following sets of ASes are built step by step using this technique:

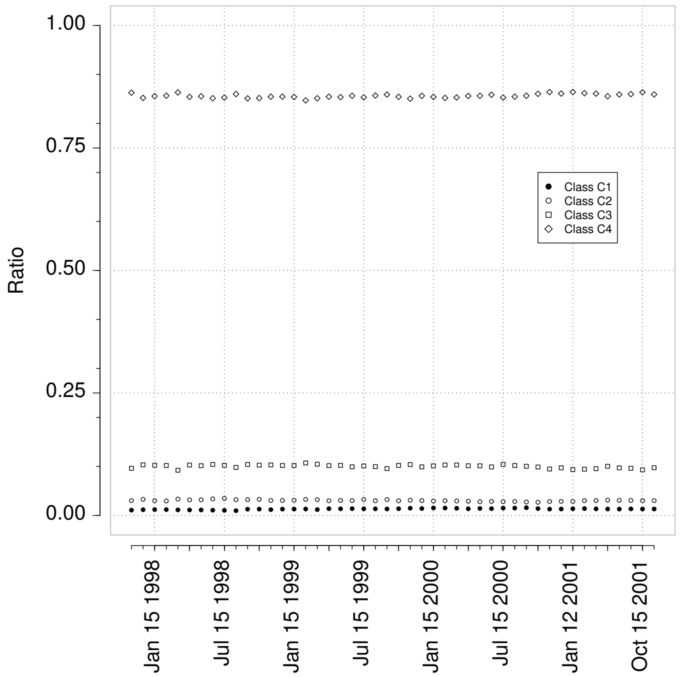

Figure 14: Classification by Govindan for Nov. 1997 till Nov. 2001

Level	Name	Number of ASes counted
0	Dense core	20
1	Transit core	129
2	Outer core	897
3	Small regional ISPs	971
4	Customers	8898

Table 3: Hierarchical classification of ASes based on data from April 2001 ([SARK02])

1. Identifying and removing nodes with outdegree 0 results in the *customer networks* (or stubs) of the AS graph.

2. By repeatedly removing nodes with outdegree 0 the *small regional ISPs* are identified.

3. The *dense core* of the Internet corresponds to the largest subset of ASes whose induced subgraph is “almost a clique” or *dense*. A directed graph of N nodes is said to be dense if each node has an indegree and an outdegree of at least $N/2$.

4. The smallest set of ASes primarily peering with each other and ASes in the dense core make up the *transit core* of the Internet.

5. The remaining nodes are the *outer core*.

This inference technique gives a hierarchical view on the Internet resulting in five levels as illustrated in Table 3.

7.2.3 Topology generators

There essentially exist two important topology generators which use a hierarchical model for generating network topologies. They will be presented in more detail in the next two subsections. Recently, yet another generator named Nbnd was proposed in [CSP02]. Nbnd is basically a modified version of Gt-Itm, producing a hierarchy of three or four levels[25], corresponding to global area networks, national area networks, MANs, and LANs respectively. Like tiers and Gt-Itm, this generator produces *spatial* graphs ([WS98]) where each node is placed on a plane and gets a pair of coordinates assigned to it.

The generator Brite also offers the possibility of using hierarchies—this generator will be looked at in Section 7.5.3.1.

7.2.3.1 tiers

The topology generator tiers, one of the first generators for network topologies, was written by M. B. Doar ([Doa96, CDZ96]) in 1996. tiers constructs 3-level hierarchies of a WAN, MANs, and LANs following the Transit-Stub model described above. The user specifies the size of each level and the number of MANs/WAN and LANs/MAN. In addition, the average number of extra edges within each level and between successive levels may be specified—Figure 15 shows the resulting network for a sample input.

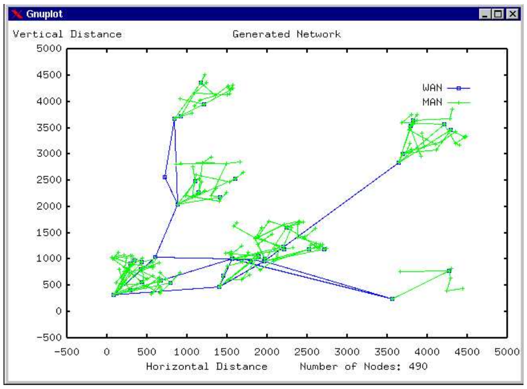

Figure 15: tiers

tiers can be used for generating network topologies in simulation projects. Some simulation tools contain conversion programs to convert output generate by tiers in the respective simulation language (e. g. the network simulator ns-2 [ns02]).

7.2.3.2 Gt-Itm

Gt-Itm[26], as proposed in [ZCB96, ZCD97] in 1996, models a three-level hierarchy corresponding to transit domains, stub domains, and LANs attached to stub nodes. This approach, called Transit-Stub model, is similar to the one chosen by tiers. Extra edges between stub domains and transit nodes are added by random selection of domains and nodes.

---

[25] Actually, Nbnd uses two different generation methods, namely Nbnd-G for global graphs, and Nbnd-N for “country” graphs.

[26] Georgia Tech Internetwork Topology Models.

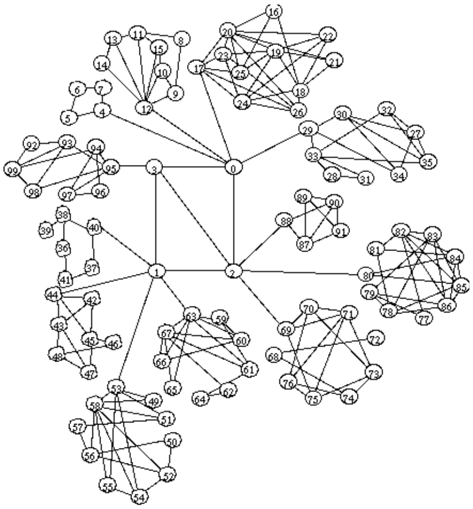

Figure 16: GT-ITM

The generator is parameterized as follows:

- T, N_t: Parameter T specifies the number of transit ASes in the network, each of which holds N_t nodes in average.
- K, N_s: On the average each transit node connects to K stub domains, and each stub domain contains N_s nodes.
- E_t, E_s: These parameters define the number of extra links between transit domains and stub domains (E_t) and among different stub domains (E_s).

GT-ITM adds several attributes to nodes and edges of the generated graph: nodes have several identifiers indicating the type of the node and the domain it belongs to (Figure 16).

7.2.4 Assessment of hierarchical models

Although hierarchical approaches seem to be promising, graphs produced by the aforementioned models do not capture all topological properties of AS graphs. In particular, it was shown that the distribution of the node degrees of topologies produced with TIERS or GT-ITM does not obey a power-law ([JCJ00, MMB00, MLMB01]).

Nevertheless, the Internet is somewhat hierarchical as the presented heuristic methods show. ISPs are often segmented into a small number of *tiers.* Usually, an ISP purchases services from a provider at a higher tier, and sells its service to customers at a lower tier. In this way ISPs can be grouped and hierarchically ordered. However, the induced hierarchy is by no means strict as many ISPs also peer with ISPs in the same tier, and more and more customer ASes are multi-homed.

7.3 Small-world graphs

Small-world graphs, introduced by D. J. Watts and S. H. Strogatz [WS98][27], have the following topological properties:

---

[27] Watts and Strogatz also call them “small-world” networks.

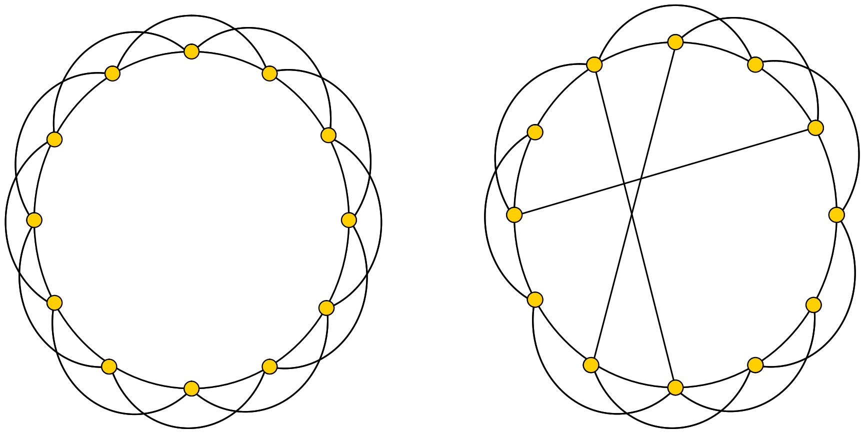

(a) 1-lattice ($n = 12$, average degree 4)

(b) After 3 random rewirings

Figure 17: Constructing a small-world graph

- They are highly clustered—their clustering coefficient γ is significantly higher than the one measured in random graphs. High clustering can be found e.g. in regular lattice graphs.
- Small-world graphs have small characteristic path lengths, like in random graphs.

There exist many prominent real-life examples of small-world graphs which have been investigated by people in different fields of research:

1. D. J. Watts examined the movie-actor network[28] (also called the "Kevin Bacon Graph") where the nodes represent the actors and edges indicate that the two actors were cast at least once in the same movie ([Wat99]). The undirected graph consists of 225,226 nodes with an average outdegree of 61 and has the characteristic path length $L = 3.65$. The clustering coefficient γ is 0.79 which is characteristic for small-world graphs.
2. So-called scientific collaboration graphs C have been the subject of research lately. The undirected graph C has "researchers" as vertices and edges $e = (u, v)$ if u and v have published a joint research paper. Of special interest has been the "Erdös Number graph" (e. g. [GI95]) where the vertices are mathematicians and one distinguished vertex p in C represents Paul Erdös. The distance from a vertex v to p is known as *v's Erdös number.* Thus, for example, co-authors of Paul Erdös have the Erdös number 1.
3. D. J. Watts also studied the "Western States Power Graph," which maps the power stations and high-voltage transmission lines that supply power in Southern California, and the graph representing the neural network of the tiny worm *Caenorhabditis elegans*. Both graphs exhibit small-world properties, like clustering and small characteristic path length.

---

[28] The adjacency matrix of the graph was generated using the Internet Movie database, `http://www.imdb.com`.

4. A. L. Barabási et al. examined in [BAJ00] the World-Wide Web represented as a directed graph and found that its average diameter amounts to only 19, i. e. two randomly chosen documents on the web are on average 19 clicks away from each other.

5. W. Aiello et al. ([ACL00]) examined so-called *call graphs* of long-distance phone calls. Every complete phone call represents an edge in the directed graph and every phone number which either received or originated a phone call defines a node in the graph.

D. J. Watts and S. H. Strogatz ([WS98]) offered in 1998 a mathematical explanation for the results of an experiment performed in the 1960s at Harvard by social psychologist Stanley Milgram. He gave letters to randomly chosen residents of Omaha, Nebraska, and asked them to deliver the letters to people in Massachusetts by passing them from one person to another. The average number of steps turned out to be about six, giving rise to the popular notion of “six degrees of separation.”

7.3.1 The β-model

Watts and Strogatz propose in their β-model a simple construction algorithm for producing small-world graphs ([Wat99]). The parameter β controls the randomness in the produced graph and allows for tuning the graph between regularity ($\beta = 0$) and disorder ($\beta = 1.0$). They start with a perfect 1-lattice (see Figure 17(a) for an example) whose edges are randomly rewired with probability β, according to the following algorithm:

1. Each vertex v_i is chosen in turn, along with the edge that connects it to its nearest neighbor in a clockwise sense (i. e. edge (v_i, v_{i+1})).

2. According to a uniform random deviate r, if $r < \beta$ the edge is deleted and rewired such that v_i is connected to another vertex v_j, which is chosen uniformly at random from the entire graph (excluding loops or multi-edges). If $r \geq \beta$, the edge is unaltered.

3. When all vertices have been considered once, the procedure is repeated for edges that connect each vertex to its next-nearest neighbor (i. e. (v_i, v_{i+2})), and so on. In total $k/2$ rounds are completed, until all edges in the graph have been considered for rewiring exactly once.

Figure 18 illustrates how the choice of parameter β affects the randomness of the produced graph. If β is set to zero, a regular 1-lattice with n vertices and average node degree k will be returned, whereas for $\beta = 1.0$ each edge of the 1-lattice will be randomly rewired resulting in a random graph without any topological order.

7.3.2 Topology generators

To the authors’ best knowledge, no publicly available topology generator has implemented this model for generating network topologies so far.

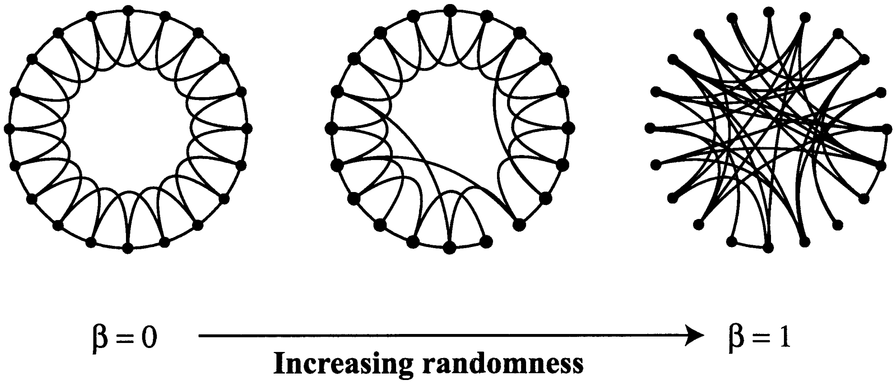

Figure 18: Schema of β-graph construction, taken from [WS98]

7.3.3 Assessment of small-world graphs

GINT provides an implementation of the β-model. We had the experience that the choice of parameter β is crucial for generating small-world networks: If β is set too small, too much randomness is added to the graph resulting in low clustering and possible disconnectedness of the graph. On the other hand, if β is set too high, the characteristic path length will be relatively large. Although the graphs produced with the β-model display topological properties typical for small-world graphs we think that the β-model is not suitable for modeling Internet-like graphs. Particularly, the degree distribution of the nodes in the produced graphs is not similar to the one measured in AS graphs.

7.4 Power-laws

Zipf's law, named after the linguist G. K. Zipf (1902–1950, [Zip49]), is the observation that often the frequency f of occurrence of some event E, as a function of the rank x when the rank is determined by the above frequency of occurrence, is a power-law function:

$$f_E \quad \propto \quad \mathrm{rank}_E{}^{\alpha}$$

where the exponent α is constant.

Zipf examined English texts and counted the occurrences of the words in it. When sorting the resulting histogram by rank, with the most frequently appearing word first, then the frequency of the appearance of a word in the text as a function of its rank is described by a power-law. If the Zipf curve is plotted on a log-log scale, it appears as a straight line with a slope of roughly -1. Zipf also noticed that the populations of cities (or population of communities) follows a similar distribution.

The history of power-law distributions goes further back in history, though. In the 19th century, Italian economist Vilfred Pareto (1848–1923) was interested in the distribution of income. Instead of asking what the ith largest income is, he asked how many people have an income greater than x. Pareto's law is given in terms of the cumulative distribution function:

$$P[X \geq x] = \left(\frac{x}{k}\right)^{-\alpha}$$

Often, Pareto's Law is interpreted as a rule of thumb stating that 20 % of a population earns 80 % of its income.

Zipf's discovery has had a great impact on later work in many different fields of research, namely:[29]

- Linguistics, particularly research on natural languages
- Study of city populations
- Bibliometrics, library science
- Finance and business
- Web access statistics and Internet traffic

7.4.1 Power-laws of the Internet

In 1999, Faloutsos et al. examined in [FFF99] several network topologies on router level and AS level and made some interesting discoveries. As many of the observed graph properties in Internet topologies follow highly skewed distributions, they can be captured by power-laws. The authors proposed three major power-laws characterizing the topology of the Internet. The first two power-laws deal with the outdegree of a node v:

Power-Law 1 (rank exponent) *The outdegree, $deg(v)$, of a node v, is proportional to the rank of the node, r_v, to the power of a constant, $\mathcal{R}$:*

$$deg(v) \quad \propto \quad r_v^{\mathcal{R}}$$

In Figure 19, this power-law is illustrated—in 19(a) the outdegrees of the nodes of the AS graph of November 15 1997 are plotted against their rank in a log-log scale. Figure 19(b) shows the rank exponents for the whole observation period. $\mathcal{R}$ lies in the interval $[-0.8232, -0.7969]$ indicating that the plots are approximated well by the linear regression.

Power-Law 2 is defined as follows:

Power-Law 2 (outdegree exponent) *The frequency f_d of an outdegree deg is proportional to the outdegree to the power of a constant, $\mathcal{O}$:*

$$f_{deg} \quad \propto \quad deg^{\mathcal{O}}$$

To get a better approximation of the data, Faloutsos et al. excluded a small percentage of nodes of higher outdegree that have a frequency of one ([FFF99]). They only consider outdegrees starting from one until they reach an outdegree that has frequency of one. In Figure 20, this power-law has been verified for the data of November 15 1997 (20(a)) and for the whole observation period (20(b)).

Interestingly enough, if all outdegrees are plotted (i. e. also nodes of higher outdegrees; compare with Figure 21), the variation of the outdegree exponent is considerably smaller

---

[29] Among other researchers Noam Chomsky (linguist), George Miller (mathematical psychologist), Herbert Simon (Nobel Prize winning economist), and Benoit Mandelbrot (computational physicist) have attempted to give reasons for the ubiquity of data obeying Zipf's Law.

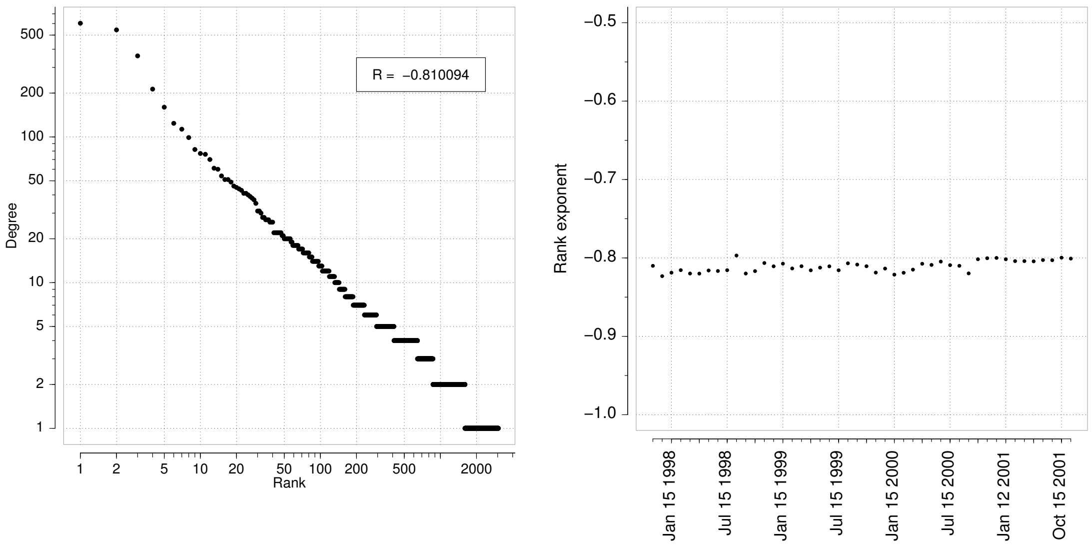

(a) Rank exponent for November 15 1997 (b) Rank exponents for observation period

Figure 19: Rank exponents of AS graphs

during the observation period (see 20(b)) than if a small percentage of nodes (about 2 %) is excluded from the plot. The outdegree exponents, though, differ significantly from the ones computed by Faloutsos et al.

Power-Law 3 (eigen exponent) *The eigenvalues λ_i of a graph are proportional to the order i to the power of a constant, $\mathcal{E}$:*

$$\lambda_i \quad \propto \quad i^{\mathcal{E}}$$

The authors plotted the first 20 eigenvalues λ_i of the adjacency matrix of the examined graphs versus its order i in log-log scale.[30] They observed an almost linear correlation with eigen exponents $\mathcal{E} \in [-0.50, -0.47]$.

In a recent article ([MP01a]), Magoni and Pansiot define four more power-laws which attempt to further describe the topology of the Internet.

7.4.2 Topology generators

Generating networks based on the degree-distribution of the nodes has been of increased interest lately, resulting in at least three topology generators. Although the construction algorithms used are different, all network topologies created with these generators have in common that the degree distribution of the nodes follows a power-law.

7.4.2.1 Inet

INET[31], described in [JCJ00], generates an AS-level representation of the Internet with similar connectivity. It generates random networks with characteristics similar to those

---

[30] The order i of the eigenvalues is implicitly given since $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$.

[31] Internet Topology Generator.

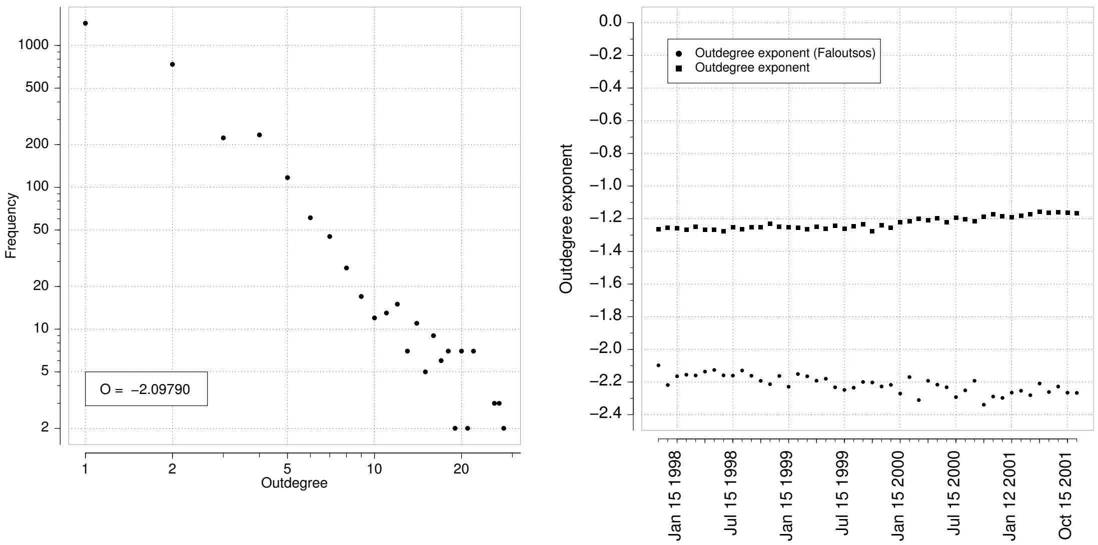

(a) Outdegree exponent for November 15 1997 (b) Outdegree exponents for observation period

Figure 20: Outdegree exponent

of the Internet from November 1997 to June 2000, and takes only two parameters:

- **N**: The total number of nodes in the produced graph.
- **k**: This parameter specifies the fraction of nodes which will be of degree 1.

INET is a typical representative of the class of degree-based graph generators. Starting from the total number of nodes N, INET computes the number of months t it would take the Internet to grow from its size in November 1997 (3037 nodes) to N, assuming exponential growth of the number of ASes. Using t, the outdegree and frequency growth can then be computed (see [JCJ00] for more detailed information). Thus, each node gets a designated degree assigned to it, and the nodes are connected to each other till their current degree reaches their designated degree.

Graphs generated in this way exhibit degree-based power-law properties well. Figure 22 shows diameter and characteristic path length of graphs generated with INET. We parameterized the generator based on the RVP dataset. Parameter N was set to the total number of vertices found in the respective real AS graph, parameter k contained the fraction of nodes of degree 1. Interestingly, the characteristic path length is well matched by INET. However, the diameters measured in the produced graphs are significantly higher and more variable than the values measured in the Internet (see Figure 22(a)).

What is more, INET can only be used for large simulation projects as all generated topologies will contain at least 3037 nodes, which is the number of ASes in the Internet in November 1997 according to the *RVP* data set.

7.4.2.2 PLRG

In [ACL00], W. Aiello et al. propose a power-law random graph model $P(\alpha, \beta)$, or (α, β)-model, where a random graph is given by a degree distribution depending on two given

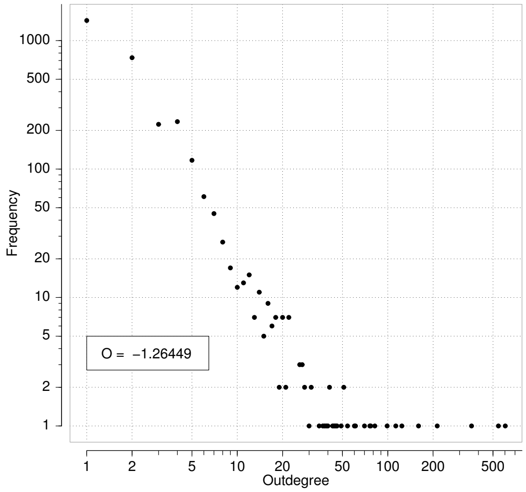

Figure 21: Outdegree exponent computed without excluding nodes

parameters α and β. Let y be the number of nodes with degree x where x and y satisfy the following condition:

$$\log y = \alpha - \beta \log x$$

This leads to the following degree distribution:

$$\left|\{v \in G | \deg(v) = x\}\right| = y = \frac{e^{\alpha}}{x^{\beta}}$$

It also follows that the maximum degree of the graph is $e^{\frac{\alpha}{\beta}}$.

The authors propose the following graph construction algorithm:

1. Form a set L containing $\deg(v)$ distinct copies of each vertex v.
2. Choose a random matching of the elements of L.
3. For two vertices u and v, the number of edges joining u and v is equal to the number of edges in the matching of L joining copies of u to copies of v.

This construction model produces multi-graphs, possibly containing loops. Comparisons with massive graphs, namely *phone call graphs*, have shown that the (α, β)-random graph model does not capture all of the random behavior of call graphs ([ACL00]). Further research has to be done to see whether this model can be used to produce graphs whose topology appears to be similar to that of the Internet.

7.4.2.3 PLOD

Another topology generator based on the distribution of node degrees was implemented by C. R. Palmer and J. G. Steffan in 2000 ([PS00]). The Power-Law Out-Degree algorithm, PLOD, uses a power-law (similar to INET) to guide the construction of the graph. In a

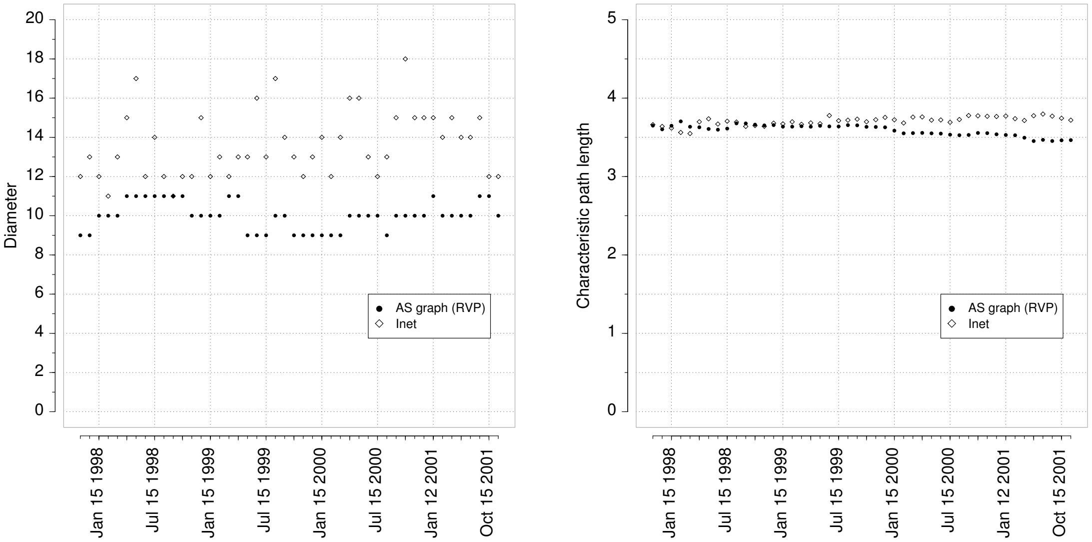

(a) Diameter (b) Characteristic path length

Figure 22: Graph invariants of INET topologies

first step, N nodes v_i are generated with designated degree $\deg(v_i) := \beta x^{-\alpha}$, where x is a uniformly distributed number between 1 and N. In a second step, pairs of nodes are repeatedly chosen by random and joined by an edge until M edges have been added to the graph. Only pairs of nodes are chosen which are not adjacent and whose current outdegrees are not yet equal to their designated degree.

PLOD needs the following parameters for generating topologies:

- **N, M**: Total number of nodes and edges in the resulting graph.
- **α, β**: These two parameters control the degree-distribution of the nodes which corresponds to a power-law.

As the resulting graph may consist of several components, a post-processing step is required which transforms the graph into a connected graph.

7.4.3 Assessment of power-law graphs

Since Internet topologies exhibit several power-law properties, power-laws are useful to assess the realism of synthetically built graphs. It is arguable, though, whether power-laws really help to advance the research on Internet topologies as they do not explain why certain topological properties appear in the Internet. Besides, power-laws can be found in many different kinds of massive graphs or networks and are not restricted to Internet graphs. Thus, they do not represent a sufficient means for describing Internet topologies.

The results of Faloutsos and Barabási were recently questioned by [TDG^{+}01]. The authors are suggesting an alternative explanation—namely, that the AS size[32] determines the degree and that the high-variability in degree is a direct consequence of the observed

---

[32] As measured by the number of routers in the AS.

high-variability in AS size. Furthermore, it was observed by Chen et al. in [CCG+02] that the discovered power-laws are only found when analyzing the RVP dataset. Extended AS maps which do not only rely on the RVP dataset (see Section 5), give a more complete picture of the Internet AS graph. The authors demonstrate that the degree distributions in these extended maps are heavy-tailed or highly-variable, but only the tails generally conform to a power-law.

7.5 Evolutionary models

Evolutionary models attempt to simulate the growth of the Internet over a certain period of time. Economical, sociological, or technological forces, which are assumed to influence the growth of the Internet, are mapped to the model. The scale-free model and its extension (often called the BA-model) are the most prominent examples for modeling the Internet using evolutionary models and are presented in the following sections.

7.5.1 Scale-free model

A. L. Barabási and R. Albert proposed in [BA99] the *scale-free model*, an evolutionary network model for producing network topologies displaying power-law properties. Many large networks have the common property that the degrees of the vertices follow a scale-free power-law distribution, e. g. in the World-Wide Web [BAJ00] or the Internet at the AS-level. In contrast to power-law models, the scale-free model does not have the parameters of a power-law as input, but relies on some dynamical process. The authors made the following two observations when analyzing graphs of the World-Wide Web:

1. *Growth:* Large networks have continuously expanded over a longer period of time by the addition of new vertices.

2. *Preferential attachment:* New vertices prefer to connect to an existing vertex of a higher degree than to one of a lower degree. That is, the probability with which a new vertex connects to an existing vertex does not obey a uniform distribution, but there is a higher probability that the new vertex will be linked to a vertex that already has a large number of connections.

Based on these observations, the scale-free model is defined as follows:

- Starting with m_0 vertices, at every timestep i a new vertex v_i with $m \le m_0$ edges is added to the graph (*growth*).

- The vertices which v_i connects to are chosen based on the probability function

$$\Pi(v) = \frac{\deg(v)}{\sum_{v_j \in V} \deg(v_j)} \qquad v \in V$$

After t timesteps the model leads to a graph $G = (V, E)$ with $|V| = m_0 + t$ vertices and $|E| = m \cdot t$ edges. As shown in [BAJ00], graphs produced in this way display power-law properties.

7.5.2 Barabási-Albert model

In [AB00], R. Albert and A. L. Barabási propose an extended model of network evolution which incorporates three basic operations: the addition of vertices, the addition of edges, and the rewiring of edges. The BA-model is an extension to the scale-free model and works as follows:

- In a first step, m_0 isolated vertices are generated.
- Similar to the scale-free model, the algorithm iterates for a time period t, and at each timestep one of the following operations is performed:
  - With probability p, $m \leq m_0$ edges are added. One endpoint of the edge is randomly chosen, the other endpoint $v \in V$ is selected with probability
    $$\Pi'(v) \quad = \quad \frac{\deg(v) + 1}{\sum_{v_j \in V} (\deg(v_j) + 1)}$$
  - With probability q, $m \leq m_0$ edges are rewired. A vertex v_i and an edge l_{ij} connected to v_i is randomly chosen. This edge is removed and replaced with a new edge l_{ik} connecting vertex v_i with v_k chosen with probability $\Pi'(v_k)$.
  - With probability $1 - p - q$, a new vertex v is added. Also, m new edges are added which connect v to m existing vertices v_i with probability $\Pi'(v_i)$.

For $p = q = 0$ the BA-model corresponds to the scale-free model. The probability function Π', expressing preferential attachment, is $\neq 0$ for all $v \in V$, in contrast to the function Π in the scale-free model. Thus, newly added or rewired edges may also connect to isolated vertices. Also note that due to the rewiring of edges, the resulting graph may be disconnected.

7.5.3 Topology generators

The topology generator BRITE implements the scale-free model and is presented in the next paragraph. However, no implementations of the Barabási-Albert model seem to be publicly available. Some confusion also exists regarding the name of the scale-free model: Sometimes, this model is also called the BA-model. In this report, the BA-model exclusively denotes the extension to the scale-free model presented in the previous section.

7.5.3.1 BRITE

The topology generator BRITE[33] proposed in [MMB00] implements the evolutionary model similar to the suggested by Barabási and Albert in the scale-free model.

The topology generator is parameterized as follows:

- **N**: The total number of nodes in the generated topology are set by this parameter.
- **HS, LS**: The nodes of the generated topology are distributed in a plane divided into HS x HS squares. Each of these high-level squares is further subdivided into LS x LS low-level squares, each of which can hold one node at maximum.

---

[33] Boston university Representative Internet Topology gEnerator. Available at `http://www.cs.bu.edu/fac/matta/software.html`.

- **NP**: This parameter controls the placement of the nodes on the plane, which can be either at random or heavy-tailed.
- **m**: Whenever a new node joins the topology, it will be connected to m nodes of the existing topology.
- **PC**: This parameter is used to activate preferential connectivity and locality as suggested by the dynamic growth model for scale-free networks.
- **IC**: When incremental growth is disabled, all nodes are first placed on the plane before edges are added to the graph. If enabled, nodes successively join the graph.

The authors suggest to have parameters *PC* and *IC* activated to produce topologies which exhibit power-law properties. ([MMB00]). BRITE also supports random graphs and can be used to produce topologies at the router-level. Using a top-down approach or bottom-up approach respectively, an AS graph can be built which also contains router nodes.

7.5.4 Assessment of scale-free graphs

The scale-free model represents an interesting evolutionary model which attempts to explain the formation of the Internet. It was shown that the degree distribution of the created nodes obeys a power-law ([BAJ00, MMB00]). However, our analysis shows that some graph invariants of the produced networks do not match those found in real Internet topologies (e. g. the RVP data set). Figure 23 shows diameter and the characteristic path length of a series of topologies created with BRITE, where parameter N corresponds to the number of nodes in the respective records of the RVP data set.

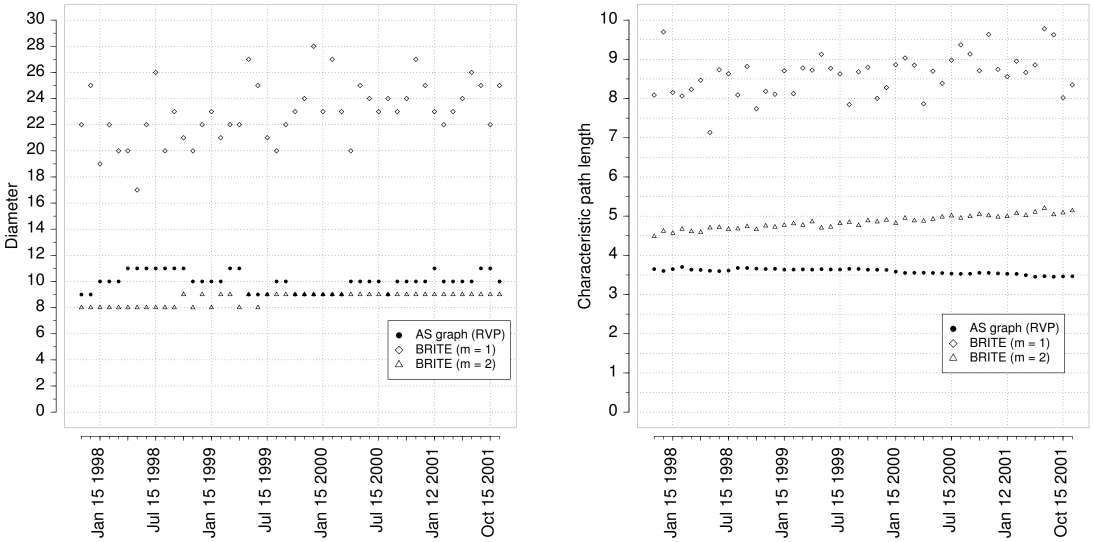

(a) Diameter (b) Characteristic path length

Figure 23: Graph invariants of BRITE topologies

Both diameter and characteristic path length are significantly higher for $m = 1$ than in the RVP data set. In order to produce graphs containing more edges than nodes, parameter

m has to be set to a value higher than 1 which results in graphs which have hardly any nodes of degree 1. In contrast, most nodes in Internet topologies are of degree 1 (see Section 6). Furthermore, the node degrees are not highly variable—the maximum node degree of topologies produced with BRITE is many magnitudes smaller (for $m = 1, 2$) than in real AS graphs. This was also observed in [CCG+02], stating that the model of Barabási-Albert does not explain the structure of the vertex degree distributions in the resulting graphs.

7.6 Spectral analysis

The spectrum of a graph G, defined as the set of eigenvalues of G[34], offers another way to analyze graphs. Eigenvalues are closely related to many other graph invariants ([FFF99]) and can be used to assess the quality of synthetically generated graphs. Although spectral graph theory has a long history, only recently Internet researchers have begun to examine the spectrum of AS graphs ([GKK01, VHE02]).

D. Vukadinović et al. investigated in [VHE02] the normalized Laplacian spectrum *nls* of Internet topology graphs. The *nls* is the set of eigenvalues of the normalized Laplacian of a graph $\mathcal{L}(G)$ which is a quadratic matrix defined on the graph ([Chu97]):

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

$\mathcal{L}(G)$ is very similar to the standard adjacency matrix of a graph, but has the advantage that all eigenvalues of $\mathcal{L}(G)$ are in the interval $[0, 2]$. The authors note that the spectra of all analyzed Internet AS-level snapshots were similar, revealing a relatively large multiplicity of eigenvalue 1. Topologies generated with INET (version 2.1) have a different *n*ls than in real AS graphs. They propose a structural classification of the nodes of an AS graph G in four different sets, P, Q, R, and I which corresponds roughly to the (single-homed) stubs, core, AS alliances, and multi-homed stubs. As most AS graphs seem to have characteristic *nls*, the spectrum can be used as a fingerprint of a graph and represents another way of evaluating synthetically generated topologies.

7.6.1 Topology generators

To the authors' knowledge no topology generator exists which produces graphs based on the spectrum.

7.6.2 Assessment of methods based on the graph spectrum

The spectrum of a graph can be used to evaluate the quality of generated topologies. However, no algorithms have been proposed yet which allow the generation of graphs which have a given spectrum. Therefore, no assessment can be made in this regard.

---

[34] Sometimes, the spectrum of a graph is defined to include also the multiplicities of the eigenvalues ([Big96]).

7.7 Other approaches

D. Magoni and J. J. Pansiot have developed the topology generator NEM [35] which provides several generation methods for Internet topologies. For instance, it contains implementations of the algorithms PLOD and PLRG. Also, the authors propose a new algorithm for generating graphs based on a random sampling of routing tables available from BGP routers. The generated graphs can be several magnitudes smaller than the graph corresponding to the original routing table. According to the authors, the produced graphs obey the power-laws introduced in this section.

Using co-called cluster graphs, B. Krishnamurthy and J. Wang propose in [KW01] another way of creating network topologies which attempt to be similar to the Internet. A cluster graph models the Internet at the level of CIDR addresses, and hence is more fine-grained than an AS graph. IP addresses having the same longest matched prefix are grouped into one cluster which is identified by the shared prefix. A cluster graph is an undirected graph with a node representing a cluster of routers and hosts which share the same prefix. An edge represents an inter-cluster connection. The authors construct cluster graphs based on BGP routing tables (*hierarchical* cluster graph), traceroute data (*traceroute-based* cluster graph), or power-laws (*synthetic* cluster graph). The topological properties of cluster graphs have not been extensively investigated yet. Also, as far as we know, no publicly available topology generator provides methods for generating cluster graphs.

In this section we have analyzed and assessed the major modeling approaches for modeling the Internet topology. Although numerous models have been proposed in the past, none of them really succeeds in producing realistic AS graphs.

8 Concluding remarks

Our report gives an overview of the current state of Internet topology research. It provides detailed theoretical background information with regards to Internet topology research and discussed several important applications of Internet topology models. The study of the Internet topology usually consists of an iterative four-stage process including the acquisition of topology data of the Internet, which is processed into a graph (usually undirected) and then analyzed. Based on these results, a model can be defined which needs to be evaluated using different sets of snapshot data. Many modeling approaches exist today which attempt to produce “realistic” topologies. They can be classified into two groups: *descriptive* models and *explanatory* models. In the last few years, the majority of the proposed models have been descriptive models, particularly degree-based approaches. Explanatory models are a lot more difficult to define in order to produce graphs which resemble Internet topologies. So far, hierarchical or scale-free approaches have not succeeded in generating graphs with the desired topological properties.

In our opinion, it is doubtful whether there will ever exist a universal modeling approach which will satisfyingly succeed in producing “realistic” Internet topologies. Al-

---

[35] Network manipulator, previously NM. The latest version of NEM (we tested version 0.95) can be retrieved from `http://clarinet.u-strasbg.fr/nem/`.

though fairly good results can be produced using topology generators based on newer descriptive models, like INET, recently discovered topological properties of AS graphs are absent in synthetically generated graphs (e. g. maximum clique size or clustering coefficient are significantly different). Too much attention has been given lately to power-laws found in AS graphs and generator methods based solely on degree distributions. Although interesting and useful for evaluating synthetically produced topologies, power-laws do not fully describe network topologies like the AS graph. This becomes apparent when examining the available topology generators which produce graphs with power-law properties (particularly Power-law 1). Major graph invariants measured in these graphs do not correspond well with those found in AS graphs, like diameter or maximum degree. Therefore, more or other input parameters are needed when generating network topologies with "realistic" topological properties. Descriptive models need to be further improved in the future so that the produced topologies also exhibit properties like having a fairly large maximum clique or a higher clustering coefficient, similar to the Internet topology. Several other factors are given below which make it considerably more difficult to find the right modeling approach:

- Available data describing the Internet topology at a certain level is incomplete and inaccurate. Therefore, statistical analysis of topology data may yield incorrect results and hence bad topology models.

- Not all major topological properties of the Internet have been investigated yet. Although prominent graph invariants of Internet graphs like size, average degree, degree distribution, diameter and the like have been extensively examined, these properties do not fully describe the topology of the Internet.

- As it is not really known how and why the topology of the Internet has evolved as it has, it is difficult to produce "realistic" graphs using explanatory models. The complex processes governing the growth of the Internet are numerous and difficult to define. In explanatory models, simple rules are used which are usually based on assumptions or on observations (e. g. preferential attachment or random rewiring in the BA-model).

- Not only the size (in terms of nodes and edges) of the Internet is changing over time, but also its structure—a fact which is often neglected when designing topology generators. Some properties of the AS graph in 1997 have significantly changed in the meantime. For instance, average node degree of the AS graphs has only slightly increased over the past four years, but the maximum degree has more than quadrupled. The ratio of single-homed stub domains (i. e nodes of degree 1) has decreased, as opposed to the ratio of multi-homed stubs which increased. Apparently, more customer networks are willing to make peering agreements with different ISPs, thus obtaining a higher resilience of their Internet connection. Probably, different models have to be used when generating topologies for different points in time.

- Many topology generators use *spatial* graphs when modeling the Internet at the AS level. This is problematic as the mapping between ASes and their geographic position is ambiguous. Since an AS is determined by its set of internal and external routing policies, ASes can geographically expand over several countries or continents. Sometimes, the geographical position of the AS headquarters are used instead ([Coo02]) although this choice of coordinates represents an strong simplification. The AS graph is actually a *relational* graph without any spatial information. This kind of information is not contained in routing tables stored in BGP routers.

What are some remaining open problems relating to Internet topology research that need to be addressed in the future? Three important areas of research are given below:

1. All topology generators available today produce simple network topologies, i. e. unlabelled AS graphs. Next-generation topology generators should be able to produce attributed graphs, where the nodes and the edges have additional attributes. For instance, a node representing an AS can contain an attribute specifying the size of an AS in terms of hosts and routers contained in the respective AS. Another attribute may define the function of a node, e. g. international backbone, national networks, regional providers, and corporate or academic networks. Edges could be classified into provider-to-customer, customer-to-provider, peer-to-peer, or sibling-to-sibling relationships. An attributed AS graph allows to create better and more realistic simulation scenarios of the Internet.

2. A lot of time has been spent on the research of the Internet topology with regards to its connectivity. Little research has been done so far regarding the reachability information in AS graphs, mainly since this kind of information is hard to get hold of. However, reachability information is not included in the current topology models. Adding semantical information to the AS graph will be of great value to simulation projects and should be the goal of future research.

3. Many metrics, laws, and graph invariants exist today which attempt to describe the Internet topology. Undoubtedly, other metrics exist which will further characterize it and which will also be useful for defining topology generators. Therefore, when new graph invariants or metrics are identified, existing topology generators need to be reevaluated and possibly adapted. It remains to be seen whether the properties that hold for the Internet topology today, will still be of importance in the future. It is possible that the structure of the Internet will significantly change so that current topology generators will have to be adjusted accordingly.

We hope to have made clear that the problem of modeling the Internet topology is a complex one, and it is not to be expected that an exact model can be found. Fortunately, it is often sufficient to have a good approximative model which can be used in large simulation projects or for the design of new distributed Internet applications. As it is the case with many problems in economy or science, suitable approximations have to be used due to the lack of complete information on the object to be modeled, and frequently they are good enough to find the expected solution.

References

[AB00] Réka Albert and Albert-László Barabási. Topology of evolving networks: local events and universality. In *Physical Review Letters*, volume 85, pages 5234–5237, 2000.

[ACL00] William Aiello, Fan Chung, and Linyuan Lu. A random graph model for massive graphs. In *Proceedings of the Thirtysecond Annual ACM Symposium on Theory of Computing*, pages 171-180, 2000.

[AJB00] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. In *Nature*, volume 406, pages 378–382, July 27 2000.

[ALPH01] Lada A. Adamic, Rajan M. Lukose, Amit R. Puniyani, and Bernardo A. Huberman. Search in power-law networks. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, October 2001.

[ANTC02] University of Oregon Advanced Network Technology Center. University of oregon route views project. `http://www.routeviews.org/`, January 2002.

[AU99] Antony Antony and Henk Uijterwaal. Routing information service r.i.s.—design note. `http://www.ripe.net/ris/Notes/ripe-200`, October 1999.

[AVG$^+$99] Cengiz Alaettinoglu, Curtis Villamizar, Elise Gerich, David Kessens, David Meyer, Tony Bates, Daniel Karrenberg, and Marten Terpstra. Routing policy specification language (rpsl). *Request for Comments 2622*, June 1999.

[BA99] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. In *Science*, volume 286, pages 509–512, October 15 1999.

[BAJ00] Albert-László Barabási, Réka Albert, and Hawoong Jeong. Scale-free characteristics of random networks: the topology of the world-wide web. In *Physica A*, volume 281, pages 69–77, 2000.

[BEF$^+$00] Lee Breslau, Deborah Estrin, Kevin Fall, Sally Floyd, John Heidemann, Ahmed Helmy, Polly Huang, Steven McCanne, Kannan Varadhan, Ya Xu, and Haobo Yu. Advances in network simulation. *IEEE Computer*, 33(5), May 2000.

[Big96] Norman Biggs. *Algebraic Graph Theory*. Cambridge University Press, 1996.

[BK73] Coen Bron and Joep Kerbosch. Algorithm 457: Finding all cliques of an uundirected graph. In *Communications of the Association for Computing Machinery*, volume 16(9), pages 575–577, 1973.

[Bol01] Béla Bollobás. *Random graphs*. Cambridge University Press, 2001.

[CCG$^+$02] Qian Chen, Hyunseok Chang, Ramesh Govindan, Sugih Jamin, Scott J. Shenker, and Walter Willinger. The origin of power laws in internet topologies revisited. In *Proceedings of IEEE Infocom 2002*, New York, NY, June 2002.

[CDZ96] Kenneth L. Calvert, Matthew B. Doar, and Ellen W. Zegura. Modeling Internet Topology. In *Proceedings of Global Internet*, November 1996.

[Chu97] Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[CNO99] James H. Cowie, David M. Nicol, and Andy T. Ogielski. Modeling the global Internet. *Computing in Science and Engineering*, January–February 1999.

[Coo02] Cooperative Association for Internet Data Analysis (CAIDA). Visualizing Internet topology at a macroscopic scale. `http://www.caida.org/analysis/topology/as_core_network/`, April 2002.

[CSP02] Xiaotao Cai, Victor T.-S. Shi, and William Perrizo. A place-based model for internet topology. In *3rd International Conference on Internet Computing (IC 2002)*, Las Vegas, Nevada, USA, June 2002.

[DK01] Martin Dodge and Rob Kitchin. *Mapping cyberspace*. Routledge, 2001.

[DL93] Matthew Doar and Ian Leslie. How bad is naïve multicast routing? In *Proceedings of IEEE Inforcom '93*, pages 82–89. IEEE, 1993.

[Doa96] Matthew B. Doar. A better model for generating test networks. In *Proceedings of Global Internet*. IEEE, November 1996.

[FFF99] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the Internet topology. In *ACM SIGCOMM '99 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 251–262, September 1999.

[FLYV93] Vince Fuller, Tony Li, Jessica Yu, and Kannan Varadhan. Classless inter-domain routing (CIDR): an address assignment and aggregation strategy. *Request for Comments 1519*, September 1993.

[FP01] Sally Floyd and Vern Paxson. Difficulties in simulating the Internet. In *IEEE/ACM Transactions on Networking*, volume 9(4), pages 392–403. IEEE/ACM, August 2001.

[Gao00] Lixin Gao. On inferring Automonous System relationships in the Internet. In *Proc. IEEE Global Internet Symposium*, November 2000.

[GI95] Jerrold W. Grossman and Patrick D. F. Ion. On a portion of the well-known collaboration graph. In *Congressus Numerantium*, volume 108, pages 129–131, 1995.

[GKK01] Kwang-Il Goh, Byungnam Kahng, and Doochul Kim. Spectra and eigenvectors of scale-free networks. *Physical Review E*, 64, October 2001.

[GR97] Ramesh Govindan and Anoop Reddy. An analysis of Internet inter-domain topology and route stability. In *Proc. IEEE INFOCOM*, April 1997.

[GT00] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for internet map discovery. In *Infocom 2000*, Tel-Aviv, Israel, March 2000.

[GY99] Jonathan Gross and Jay Yellen. *Graph Theory and Its Applications*. CRC Press, Inc, January 1999.

[HFMC01] Bradley Huffaker, Marina Fomenkov, David Moore, and Kimberly C. Claffy. Macroscopic analyses of the infrastructure: Measurement and visualization of Internet connectivity and performance. In *PAM2001 – A workshop on Passive and Active Measurements*, Amsterdam, April 23–24 2001. CAIDA/SDSC/UCSD.

[HN92] Steve Hotz and Ron Nagamati. A tool for generating network topology and policy for protocol simulation purposes. Technical report, Computer Science Department, University of Southern California, 1992.

[Hus99a] Geoff Huston. Interconnection, Peering and Settlements—Part I. *The Internet Protocol Journal*, 2(1):2–16, March 1999.

[Hus99b] Geoff Huston. Interconnection, Peering and Settlements—Part II. *The Internet Protocol Journal*, 2(2):2–23, June 1999.

[Hus01] Geoff Huston. Analyzing the Internet's BGP routing table. *The Internet Protocol Journal*, 4(1):2–15, March 2001.

[JCJ00] Cheng Jin, Qian Chen, and Sugih Jamin. Inet topology generator. Technical Report CSE-TR-433-00, EECS Department, University of Michigan, 2000.

[KW01] Balachander Krishnamurthy and Jia Wang. Topology modeling via cluster graphs. *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, 2001.

[LMJ99] Craig Labovitz, G. Robert Malan, and Farnam Jahanian. Origins of Internet Routing Instability. In *Proceedings of INFOCOM '99*, New York, NY, 1999.

[MBB00] Tony McGregor, Hans-Werner Braun, and Jeff Brown. The NLANR Network Analysis Infrastructure. In *IEEE Communications*, volume 38, pages 122–128. IEEE, May 2000.

[MLMB01] Alberto Medina, Anukool Lakhina, Ibrahim Matta, and John Byers. BRITE: An approach to universal topology generation. In *Proceedings of IEEE MASCOTS '01 (Tools track)*, Cincinatti, August 2001.

[MMB00] Alberto Medina, Ibrahim Matta, and John Byers. On the origin of power laws in Internet topology. *ACM SIGCOMM Computer Communication Review*, 30(2), April 2000.

[MP01a] Damien Magoni and Jean-Jacques Pansiot. Analysis of the Autonomous System network topology. *ACM Computer Communication Review*, 31(3):26–37, July 2001.

[MP01b] Damien Magoni and Jean-Jacques Pansiot. Influence of network topology on protocol simulation. In *International Conference on Networking (ICN) 2001*, pages 762–770. LNCS, 2001.

[Nat02] National Laboratory for Applied Network Research (NLANR). Network analysis infrastructure. `http://moat.nlanr.net/`, 2002.

[Nau00] John Naughton. *A Brief History of the Future*. Phoenix, 2000.

[ns02] UCB/LBNL/VINT Network Simulator NS. http://www.isi.edu/nsnam/ns/, 2002.

[PL00] Kihong Park and Heejo Lee. A proactive approach to distributed DoS attack prevention using route-based packet filtering. Technical Report CSD-TR-00-017, Purdue University, Dept. of Computer Sciences, December 2000.

[PS00] Christopher R. Palmer and J. Gregory Steffan. Generating network topologies that obey power laws. In *IEEE Globecom 2000*, San Francisco, November 2000.

[RL95] Yakov Rekhter and Tony Li. A border gateway protocol 4 (bgp-4). *Request for Comments 1771*, March 1995.

[RTY$^+$00] Pavlin Radoslavov, Hongsuda Tangmunarunkit, Haobo Yu, Ramesh Govindan, Scott Shenker, and Deborah Estrin. On characterizing network topologies and analyzing their impact on protocol design. Technical Report CS-TR-00-731, USC/Information Sciences Institute, March 2000.

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

[Ste99] John W. Stewart. *BGP4 : Inter-Domain Routing in the Internet*. The network basics series. Addison-Wesley, 1999.

[TDG$^+$01] Hongsuda Tangmunarunkit, John Doyle, Ramesh Govindan, Sugih Jamin, Scott Shenker, and Walter Willinger. Does AS size determine degree in AS topology? *ACM SIGCOMM Computer Communication Review*, October 2001.

[TGSE01] Hongsuda Tangmunarunkit, Ramesh Govindan, Scott Shenker, and Deborah Estrin. The impact of routing policy on Internet paths. In *Infocom 2001*, Alaska, USA, April 2001.

[VHE02] Danica Vukadinovic, Polly Huang, and Thomas Erlebach. On the spectrum and structure of internet topology graphs. In *Innovative Internet Computing Systems*, Ostseebad Kühlungsborn, Germany, June 2002.

[Wat99] Duncan J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, 1999.

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

[WE94] Liming Wei and Deborah Estrin. The trade-offs of multicast trees and algorithms. In *Proceedings of ICCCN '94*, 1994.

[WS98] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of “small-world” networks. In *Nature*, volume 393, pages 440–442, June 4 1998.

[Zak02] Robert Hobbes Zakon. Hobbes' Internet Timeline FAQ. http://www.zakon.org/robert/internet/timeline/, May 2002.

[ZCB96] Ellen W. Zegura, K. Calvert, and S. Bhattacharjee. How to model an Internetwork. In *Proceedings of IEEE Infocom*. IEEE, 1996.

[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*, 5(6):770–783, December 1997.

[Zip49] George Kingsley Zipf. *Human Behaviour and the Principle of Least Effort: an Introduction to Human Ecology*. Addison-Wesley, 1949.