BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Thesis

A SIZE-BASED APPROACH TO AUTONOMOUS SYSTEM TOPOLOGY MODELING

by

MARWAN FAYED

Dipl., Sheridan College, 1996
B.A., Laurentian University, 2000

Submitted in partial fulfillment of the

requirements for the degree of

Master of Arts

2003

Approved by

First Reader

John W. Byers, Ph.D.
Assistant Professor of Computer Science

Second Reader

Mark Crovella, Ph.D.
Associate Professor of Computer Science

Second Reader

Ibrahim Matta, Ph.D.
Associate Professor of Computer Science

Acknowledgments

I would first like to thank my co-advisors, John Byers and Mark Crovella. This work would not have been possible without their direction, guidance, patience, and enthusiasm.

I would further like to thank the co-authors in the publications related to this thesis work for their contributions and explanations. My gratitude to Ibrahim Matta for serving on the review committee and Azer Bestavros for his honesty.

Lastly, I wish to thank my mom, dad, and nora for their encouragement, as well as Emily for her undying patience, enthusiasm, and encouragement. I know, Emily, that I was not around as much as you would have liked.

A SIZE-BASED APPROACH TO AUTONOMOUS SYSTEM TOPOLOGY MODELING

MARWAN FAYED

Boston University Graduate School of Arts and Sciences, 2003

Major Professors: John W. Byers, Assistant Professor of Computer Science

ABSTRACT

Recent studies note that numerous structural properties in the Internet's autonomous system (AS) graph follow highly variable distributions. The most studied of these is the AS degree (number of peering links) distribution, it is a crucial property in the understanding of the Internet's evolution and structure. The dominant model for this phenomenon is the Barabási-Albert (B-A) model. Its central feature, preferential connectivity, requires each AS to have global knowledge. This assumption of global knowledge may produce the underlying properties observed in the AS topology, but cannot explain the reasons they emerge. This thesis explores a more general explanation for the highly variable distributions observed, one where preferential connectivity and global knowledge are absent. Specifically, we explore the relationship between AS size and degree distributions and the processes by which they emerge.

We present two size models: one incorporates only the growth of hosts and ASes, and a second extends that model to include mergers of ASes. Our models are motivated by measurements of relatively unexplored data sources, as well as from novel heuristics applied to examinations of collections of BGP tables. We show that these models yield a size distribution exhibiting a power-law tail. Furthermore, in such a model, if an AS's link formation is roughly proportional to its size, then AS degree will also show high variability. We instantiate such a model with empirically derived estimates of growth rates and show that the resulting degree distribution is in good agreement with that of real AS graphs.

Contents

Acknowledgments **iii**

Abstract **iv**

Table of Contents **v**

List of Tables **vii**

List of Figures **viii**

List of Abbreviations **ix**

1 Introduction **1**

2 Related Work **7**

3 Background **10**

3.1 Public Data Sources 11

3.1.1 Registry Summaries 11

3.1.2 BGP Tables 12

3.1.3 Host History 14

3.2 Building an Intuition 14

3.2.1 First Attempts 14

3.2.2 Using Host Counts 17

4 Constructing a Simple Model **19**

4.1 A Simple Model and its Analysis 20

4.2 Estimating Growth Rates 22

4.2.1 Analytical Validation 24

4.2.2 Empirical Validation 25

5 A More Complete Model **26**

5.1 Modelling AS Mergers 26

5.2 Implications of the Model 27

5.2.1 Constant Kernel 28

5.3 Estimating the Rate of Coalescence 29

5.3.1 Validation of the Merger Model 31

6 AS Degree Formation **33**

7 Validation **35**

7.1 Empirical Data Sources 35

7.1.1 BGP Adjacency 36

7.1.2 AS Overlays 36

7.1.3 Summary of Empirical Observations 37

7.2 Constant Connectivity Models 38

7.3 Size-Based Connectivity Models 39

8 Conclusion **42**

Bibliography **45**

List of Tables

7.1 Summary of Data Sources 37
7.2 Conductance of ASes 40

List of Figures

3.1 A Typical Summary of Allocations 12
3.2 Sample (Condensed) Entries in a BGP Table, with Legend 13
3.3 Cumulative Allocations of a) AS numbers and b) blocks of IP addresses 15
3.4 Ratio of IP addresses per Autonomous System over Time 16
3.5 Ratio of Hosts (or live IPs) over IP space 17
3.6 Ratio of Hosts (or live IPs) per AS, $\frac{Hosts}{AS}$ 18
3.7 Growth of AS wrt Appearance of a Host ("host time") 18

4.1 Growth in the number of a) ASes and b) Internet hosts. 19
4.2 Left a) AS Growth, Right b) Host Growth. 23
4.3 Number of Hosts vs. Number of ASes 24
4.4 Simple Model Predictions vs. Measurements 25

5.1 AS Reappearance Time. Left: 2000, Right: 2001. 30
5.2 Rate at which ASes Merged in the 2001 Dataset. 31
5.3 Coalescence Model Predictions vs. Measurements 32

7.1 Degree Distributions Inferred from 4 Sources. 35
7.2 Predicted Degree Distribution where $x = 0.10$ 39
7.3 Diminishing Probability where $x = 0.20$. 41

List of Abbreviations

AS	Autonomous System
ISP	Internet Service Provider
BGP	Border Gateway Protocol
B-A	Barabàsi-Albert
IP	Internet Protocol
RIPE	Réseaux IP Européens Network Coordination Centre
APNIC	Asia Pacific Network Information Centre
ARIN	American Registry for Internet Numbers
IDS	Internet Domain Survey
DNS	Domain Name System
NLANR	National Laboratory for Applied Network Research
PCH	Packet Clearing House

Chapter 1

Introduction

Many aspects of the Internet's structure are relatively unknown. These gaps in our knowledge pose problems when attempting to construct representative network topologies for simulation and modeling. In addition, filling these gaps may shed light on the forces behind the Internet's growth and the ways in which the network may fail.

There are two widely studied views of the Internet corresponding to different hierarchical levels of aggregation. The router-level topology, where we consider only those machines that are physically responsible for directing traffic between source and destination, offers the finer-grained resolution of the Internet. Another view is obtained when looking at the granularity of autonomous systems (or ASes) that are responsible for the administration of a collection of these routers. For example, an average university campus might administer five or ten routers visibile on the Internet, while a large ISP (Internet Service Provider) such as Sprint or UUNet are responsible for the operation of hundreds or thousands of routers scattered geographically. Of course every router must be associated with some AS. However, at the AS level of abstraction router information is not available. In fact, most information regarding the internal structure of the AS is hidden in this view.

We can represent these views using labelled graphs. In the AS-level graph vertices represent ASes and edges represent AS-AS peering relationships, those links established between ASes for the sake of global reachability. At the router-level, graph nodes are routers and edges are network links. A node's label corresponds to the autonomous system to which the node belongs. These abstractions provide an attractive target for Internet topology generation efforts, because they make possible significant improvements in network simulation. Properly labelled graphs would allow meaningful simulation of traffic flow patterns, which are influenced strongly by interdomain routing policies. Furthermore, accurate AS labelling would allow realistic simulations of the BGP system, which is of considerable current interest.

Unfortunately, a number of gaps in our understanding prevent the construction of such labelled graphs. Principal among these is the current lack of existence of a model for the evolution of the interdomain system. Such a model should be able to answer the questions: By what processes do new ASes arise over time? How do ASes grow? By what processes do ASes merge? What determines the interconnection processes between ASes? In this thesis we seek answers to these questions. To do so we use a variety of Internet measurements, supplemented with insight from analytic models.

A particularly surprising aspect of these graphs is the prevalence of highly variable distributions [16, 23], the most well-studied of which is vertex degree. In discussing properties of the AS graph, it is useful to draw a distinction between *high variability* and *power-law tails.* High variability is a qualitative notion, referring to a probability distribution showing non-negligible values over a wide range of scales (typically at least three orders of magnitude). On the other hand, a distribution $p(\cdot)$ with power-law tails has the formal property that:

$$p(x) \sim x^{-\alpha}$$

with $\alpha > 0$, and where $a(x) \sim b(x)$ means that $\lim_{x\to\infty} a(x)/b(x) = c$.

Some authors have argued that AS vertex degree is well modeled as having power-law tails [16, 23]. Others have suggested that vertex degree does not clearly exhibit power-law tails, although it is highly variable [12]. Since such highly-variable distributions do not arise in simple random graphs, and since power-law tails do provide a simple (albeit crude) approximation for the behavior of the true distribution, a number of papers have proposed mechanisms (more complicated than purely random connection) that may give rise to power-law degree distributions in graphs [7, 22, 21].

The most prominent model attempting to explain the emergence of power-law degree distributions is the Barabási-Albert model (or B-A model) [7, 2]. In fact, it has been considered in a number of papers as a model for AS graphs [3, 10, 29, 26, 35]. The B-A model assumes the network is formed through incremental addition of nodes. In the simplest form of the model, a new node forms a connection to an existing node with probability proportional to the existing node's degree. This preferential connectivity leads to a "rich get richer" phenomenon in which high degree nodes tend to increase in degree faster than low degree nodes. Such models have two drawbacks. First, they require global knowledge of the connectivity of all other nodes and second, there is no allowance for influences other than proportion of degree when choosing neighbouring ASes. While they might be successful in reproducing observed properties of Internet graphs, they unlikely to be representative of the processes responsible for those properties.

In this thesis we present work published in [18, 17] and examine whether explanations more general than the B-A model may suffice to explain *highly variable* degree distributions in the AS graph. In trying to find such explanations, we look to the evolution of ASes themselves. We are motivated by two observations. First, the authors in [14] point out that AS degree is strongly correlated with AS size (measured in number of nodes) — and that AS size also shows a highly variable distribution. Second, we observe that during the last 10 years or so, the Internet has undergone exponential growth in both number of nodes and number of ASes. Under such conditions, we show here that highly variable AS sizes (and, presumably as a consequence, highly variable AS degrees) may readily arise due to exponential growth alone.

We explore these observations in this thesis by first constructing a simple growth model for ASes. Our model makes three assumptions: (1) exponential growth in the number of hosts in the network; (2) exponential growth in the number of ASes in the network. In this model, (the total number of hosts) and N (the total number of ASes) are described by the simple linear growth equations $dN/dt = qN$ and $dM/dt = pM + qN$, where q and p are the growth parameters. We show that in the asymptotic time limit, this model leads to a stationary size distribution with power-law tails.

The principal validation of this model is to check whether its predicted size distribution agrees with empirical measurements of AS size distributions. For this purpose we use two large router inventories, from the Mercator and Skitter projects, and map each router to its associated AS. The resulting size distributions of ASes are found to have long (though not clearly power-law) tails. Agreement between the AS size distributions predicted by the growth model and the data is a good first order approximation, but there are noticeable discrepancies. For example, the tail of the model distribution is in general agreement with data, but it strictly follows a power-law, while the empirical data shows some deviation from a power-law tail.

However, in the Internet there are events other than AS growth which occur. We hypothesize that an important factor our simple model omits is the merger of ASes. Statistical mechanics shows that highly variable size distributions can also result from coalescence processes (as in the formation of raindrops or polymer aggregates [24]). To understand the merger process and estimate a corresponding rate, we develop a heuristic and apply it to examinations of BGP table collections over two one-year periods. Adding mergers to our growth model complicates the analysis considerably. Currently only the most tractable version of this model has been solved, in which mergers occur with a rate proportional to the number of ASes present, in a manner independent of the sizes of the ASes merging. In this version the new rate of growth in the number of ASes is described by $dN/dt = (q - r)N$, where r is the rate of coalescence.

This merger model exhibits improved agreement with data with respect to small to medium sized ASes, but predicts large AS sizes less well, compared to the pure growth model. More importantly, it points to the need for analysis and detection of more realistic merger processes (such as those that account for the relative sizes of the ASes being merged).

We then return to study the impact of size on node degree and inter-AS connectivity. This extension makes a third assumption in addition to those above, that there exists (3) an approximately proportional relationship between AS size and degree [14]. The resulting model shows that highly variable AS degrees may easily arise without preferential connectivity, and in fact without any global knowledge of network state by individual ASes. Indeed, in our model, the methods by which ASes select peering partners can remain completely unspecified.

We present a simple algorithm which constructs AS-AS peering links as the ASes grow in size, over time. We show that, in an environment where the growth rate of hosts and ASes is exponential, random selection process suffices to produce a highly variable degree distribution among ASes. To demonstrate the flexibility of this model an example is presented where we specify that the selection process consider the size of the AS when deciding whether to insert new links.

We conclude that, for topology generation, it is not necessary to incorporate preferential connectivity in order to generate highly variable AS degree distributions. This leaves the door open for more practically justified bases for forming inter-AS links, *e.g.,* based on economic and geographical considerations.

This thesis provides a growth model to connect ASes. This model for the AS graph is more general than the B-A model, and is based on empirical observations of Internet growth dynamics. It allows for inter-AS connections to be formed in a way that need not be based on AS degree, losing. We show that this model yields highly-variable degree distributions, and that its outputs agree well with empirical measurements of AS graph degree distribution.

In addition, we show that growth based models are a good first step to understanding the evolution of the sizes of ASes. Our model leads naturally to a method for labelled network topology generation in which the topology grows incrementally, and as nodes are added, new ASes arise, and existing ASes grow and merge.

Chapter 2

Related Work

Until recently, Internet topologies have been generated using random and hierarchical models. Among the more significant of these is work due to Calvert *et al.* in [11]. That paper proposes generating smaller domain-like networks and connecting them together to create a hierarchical structure whose properties are specified by input parameters. The goal in this work was to emulate the types of relationships that exist on the Internet. Unfortunately, these random and hierarchical approaches fail to capture many significant attributes of Internet topology as well as the power-law models [35, 26] discussed below.

Since attention was drawn to power-laws in Internet topologies by [16], modeling efforts have shifted to reproducing these power-law properties. The most notable effort in this direction has been the Barabási-Albert *preferential attachment* model [7]. In this model, the network is formed through incremental addition of nodes. The model's key assumption is that a new node forms connections to existing nodes based the existing nodes' degrees. The probability that a new node will connect to an existing node i is proportional to $\Pi(i) = k_i/\Sigma_j k_j$, where k_i is the degree of node i. The resulting rate at which nodes acquire new edges is given by $\delta k_i/\delta t = k_i/2t$, where

t is the time elapsed from the start of the process. The resulting degree distribution exhibits a power-law tail, with a fixed exponent of $\alpha = -3$.

Later work has built upon and extended the B-A model. The same authors in [3] extended the model to allow re-wiring, in which edges may also be deleted or moved at each timestep; this allows the exponent in the power law relationship to vary. The work in [29] investigates the case where only a subset of all nodes in the network are available to choose from. With only slight modifications to the B-A model they show that a power-law degree distribution emerges. Additionally, a "generalized linear preference" model is proposed in [10] which better matches the clustering behavior and path lengths of empirical Internet measurements. These extensions have improved the flexibility of the B-A model, albeit with a corresponding increase in complexity.

The generation of power-laws through random graph models has also received considerable recent attention. An overview of existing models appears in [1], along with a single set of models which generalizes them all. In this family of models, nodes are periodically added to the graph with some probability and are initially assigned an in-weight and out-weight of 1. At each timestep, t, with some fixed probability, a new directed edge is created between nodes i and j. The probability of selecting an edge from i to j is in proportion to i's out-weight and j's in-weight, respectively. Then, the out-weight of i and the in-weight of j are increased by 1; hence, at every timestep the total in-weight (or out-weight) in the system is exactly t. This general method can generate graphs with arbitrary degree distributions, but are not proposed as realistic models for the dynamics of Internet growth.

In contrast to the approaches above which focus on reproducing statistical properties, another family of models explores the implications of optimization goals on network evolution. One such model has been suggested in [15]; it assumes that nodes arrive uniformly at random within some Euclidean space, and the newly created edges attempt to balance the distance d from its new neighbour with the desire to minimize the average number of hops h to other nodes. A new node i forms an edge to j by minimizing the weighted sum $\gamma \cdot d_{ij} + h_j$. The resulting degree distribution exhibits a power-law tail. A second optimization-based model is described in [4]; this document explores a similar heuristic but at the ISP level.

The investigation in [14] evaluates the merits of the B-A model and its applicability to the Internet. The authors conclude that, while the B-A family of models do succeed in producing power-laws, the model itself is not representative of the dynamics that drive Internet evolution: its growth processes (preferential connectivity) do not match those observed in the Internet, nor does the requisite global knowledge assumption hold. Also, they present evidence to suggest that AS-level degree distribution is not a pure power-law, though still highly variable. This conclusion is drawn when attempting to build a more complete Internet graph using multiple sources. Based on these observations, together with evidence in [34] which links degree to size, [14] suggests that other (perhaps simpler) mechanisms decide the evolution of the Internet.

The work in this thesis shows that preferential connectivity, or indeed any dependence on degree in making connection decisions, is not necessary in order for power-law degree distributions to emerge. Furthermore, ours is the first model that models highly variable degree distributions as well as the size and growth of autonomous systems themselves.

Chapter 3

Background

Before beginning a discussion of AS size distributions we provide a detailed description of the data used (that we considered using). We also describe the initial steps taken to develop an understanding of AS growth and behaviour. The intuition we seek should help us to construct a model to describe AS growth as seen in the Internet.

In order to track any entity in any study, some form of identification is required. When seeking services, in many countries people are most easily identified by a unique identification number (eg. social security in the U.S.A.). Moreover, we also need to be able to route information to individuals using a mailing address. The Internet's analogue to I.D. numbers and mailing address are *routing numbers* such as IP or AS numbers. Internet routing numbers are globally unique and so they can also be used to identify entities as they appear online, move, or disappear.

For the purpose of this discussion, we draw a distinction between the following terms:

- **IP**: We use 'IP', 'IP blocks', 'IP space', and 'IP prefix' interchangeably to refer to sets of contiguous IP addresses that are allocated to registrants on the

Internet (Section 3.1.1). These are allocated but not necessarily in use.

- **Host**: This term is used to refer to IPs that are actually in use on the Internet. Typically there is a 1:1 mapping of IP to physical machine, but this is not always the case[1]. We sometimes refer to hosts as "*live IPs*" to reflect the actual use of an IP address.

Following a comprehensive search for and scrutiny of data sources, we determined the most useful and accurate sources of information to be Internet Registrations summaries [5, 6, 31], BGP tables, and a history of the number of hosts in the Internet (if available).

We first discuss the data chosen and summarize its shortcomings. Then we briefly present our initial attempts to gain an intuition, with efforts that followed in subsequent chapters.

3.1 Public Data Sources

3.1.1 Registry Summaries

Three registries are responsible for allocating globally unique routing numbers (IP and AS numbers) worldwide. The agency RIPE is responsible for Europe and super-Sarahan Africa, APNIC for South-East Asia, and ARIN for the remainder of the world.

Each of the registries keeps a history of all allocated routing numbers that is publicly available [6, 5, 31]. This summary of allocations lacks any ownership information, as can be seen in figure 3.1. Each entry is delimited by a vertical bar and can be decoded according to the legend below.

[1]For our purposes we say that multiple IPs assigned to host represents multiple hosts

```
arin|CA|ipv4|216.254.128.0|24576|20000315|allocated
arin|US|ipv4|216.255.0.0|16384|20010416|allocated
arin|US|ipv4|216.255.64.0|8192|20010416|allocated
arin|US|ipv4|216.255.128.0|8192|20010430|allocated
arin|US|ipv4|216.255.192.0|4096|20010501|allocated
arin|US|ipv4|216.255.224.0|4096|20010501|allocated
arin|AU|ipv4|219.0.0.0|16777216|20011015|allocated
arin|AU|ipv4|220.0.0.0|16777216|20011203|allocated
iana|US|ipv4|224.0.0.0|268435456|19910522|assigned
arin|CH|asn|593|1|19900401|allocated
arin|DE|asn|575|1|19900401|allocated
arin|US|asn|257|1|19881122|allocated
arin|US|asn|296|1|19890608|allocated
arin|US|asn|149|1|19880104|allocated
arin|US|asn|85|1|19860910|allocated
arin|US|asn|133|1|19870919|allocated
arin|US|asn|713|1|19900823|allocated
arin|US|asn|218|1|19880729|allocated
arin|US|asn|522|1|19890825|allocated
```

Figure 3.1: A Typical Summary of Allocations

Registrar	Country Code	Type	First Routing Number	Range	Date	Not Applicable

The list is maintained by hand and is susceptible to human error. Nevertheless, it is the most accurate available record of AS allocations. These summaries provide valuable information: namely, we can trace (a) AS number assignments, (b) assignments of IP blocks, and (c) the date on which an assignment was made.

3.1.2 BGP Tables

In order to communicate amongst themselves, ASes need to know about each other. They need to know which ASes exist and how to route information to them. Routing

```
BGP table version is 3626864, local router ID is 198.32.162.100
Status codes: s suppressed, d damped, h history, * valid, > best,
i - internal Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Path
*  3.0.0.0          195.211.29.254      5409 6667 209 701 80 i
*                   167.142.3.6         5056 701 80 i
*  4.0.0.0          167.142.3.6         5056 1 e
*                   204.212.44.131      234 2914 1 i
*  6.0.0.0          167.142.3.6         5056 7018 7170 1455 i
*                   195.211.29.254      5409 6667 209 7170 1455 i
*  9.2.0.0/16       195.211.29.254      5409 6667 209 701 i
*>                  134.55.20.229       293 701 i
*  9.3.4.0/24       195.211.29.254      5409 6667 209 3561 1221 ?
*                   134.55.20.229       293 1 16779 1221 ?
```

Figure 3.2: Sample (Condensed) Entries in a BGP Table, with Legend

and reachability information is passed around using BGP (**B**order **G**ateway **P**rotocol) messages. As BGP messages travel through the network, routers build and store every valid path. Best paths are stored and chosen according to some policy.

Referring to the sample BGP table entries in figure 3.2, the information provided that is relevant for our purposes consists of the "Network" and "Path" columns. The *Network* column is the destination represented by longest matching IP prefix, while the *Path* column is a series of numbers. This series of numbers represents the path in terms of AS numbers, from immediate neighbour to final destination.

These tuples enable us to infer details such as average path length, coalescence, and link creation at the AS level, among others. Unfortunately, BGP is very volatile [9] and one must determine if snapshot information suffices, or if long-term averages are necessary.

3.1.3 Host History

A record of the number of hosts in use on the Internet is difficult to obtain. The reasons for this are two-fold. First, the Internet is now owned and managed by private entities unwilling to divulge details of their infrastructure. Second, even if it were possible, no central record of host activity exists. Hence, studies to determine the Internet's size cannot be validated, though we can accept them as useful estimates.

Only two known projects have attempted to determine the growth of the Internet with any success: Telcordia's Netsizer [30] is no longer available, hence we look to Internet Software Consortium's "Internet Domain Survey" [20]. IDS host counts are publicly available and have been produced semi-annually from January 1995, and quarterly from 1991-1995. Their counts are derived from a reversal of the Domain Name System's lookup process, the result of which is an approximate count of IP addresses in use. It is the best available record of which we are aware.

3.2 Building an Intuition

3.2.1 First Attempts

This step, for the most part, consisted of data collection and subsequent reorganization of the data into a usable format. What results are the plots in Figure 3.3. These plots show the cumulative allocations of routing numbers over time on different time scales. It should be noted that the (near) vertical line appearing in IP allocations around 1991-92 in Figure 3.3b) should be ignored as it represents an anomolous allocation (specifically the assignment of some 250 million multicast addresses).

But what can we determine from these plots? Perhaps there is some relationship that is visible. Though we can see how these numbers have been handed out, it is

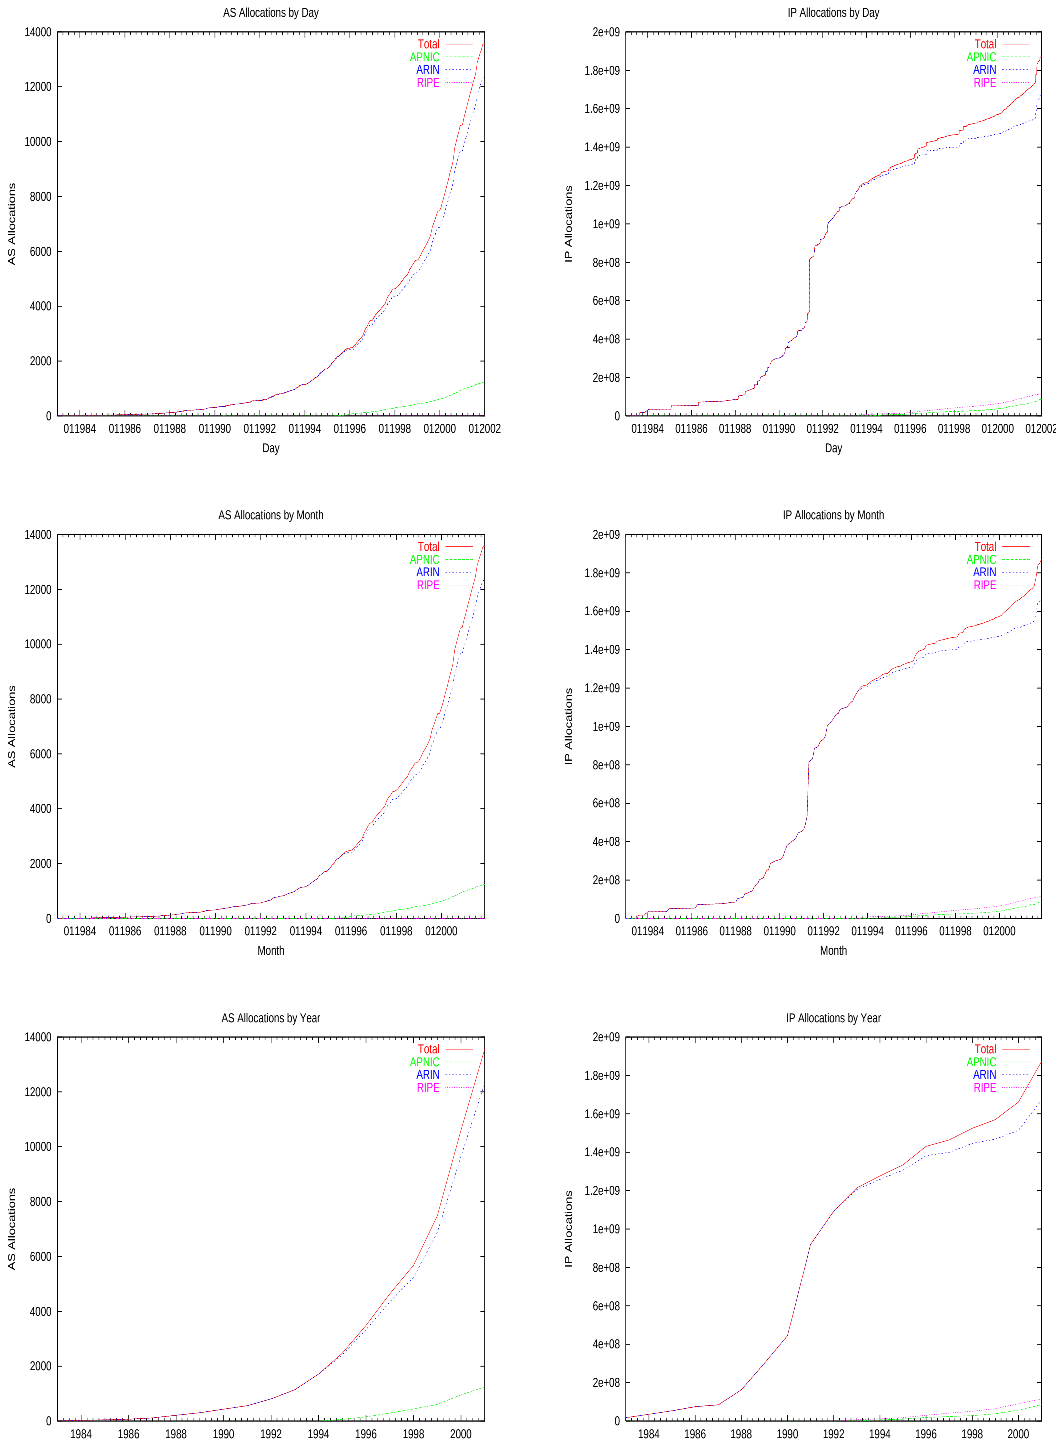

a) AS Allocations **b) IP Allocations**

Figure 3.3: Cumulative Allocations of a) AS numbers and b) blocks of IP addresses

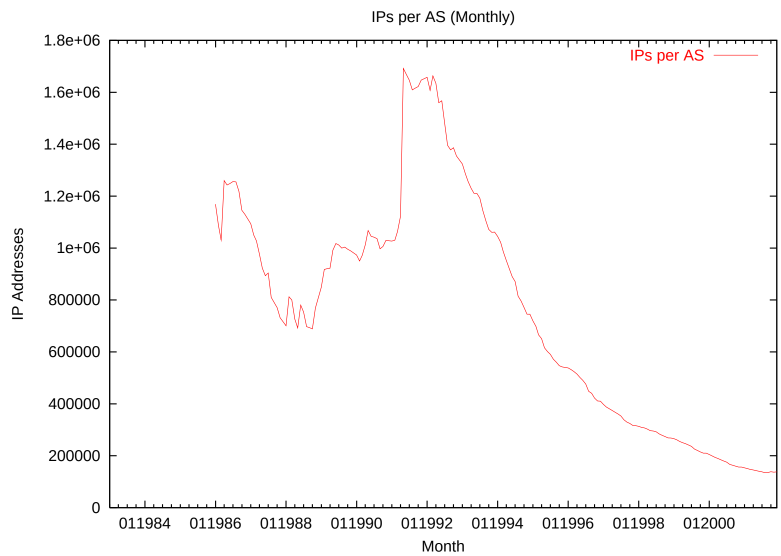

Figure 3.4: Ratio of IP addresses per Autonomous System over Time

difficult to draw any conclusions relating the number of ASes with the number of hosts.

To better understand, we next plotted the ratio of IP addresses per AS. The result is Figure 3.4. Again, the spike that appears about 1991 should be ignored for the same same anomolous allocation mentioned above.

While it may seem that there is some relationship to be inferred, we conclude that the use of IP allocations to determine host growth is misleading. There are fewer IPs handed to ASes over time simply because more IP addresses are initially allocated than are necessary. This is to accomodate growth in a network without having to repeatedly request more IP space from the registries. Also, routing table sizes can be reduced by advertising larger contiguous blocks of IP addresses.

These first attempts only reinforced a belief that IP 'space' could not be used to accurately determine the information we desire. A true host count surely would yield different results, and surely a host count did just that.

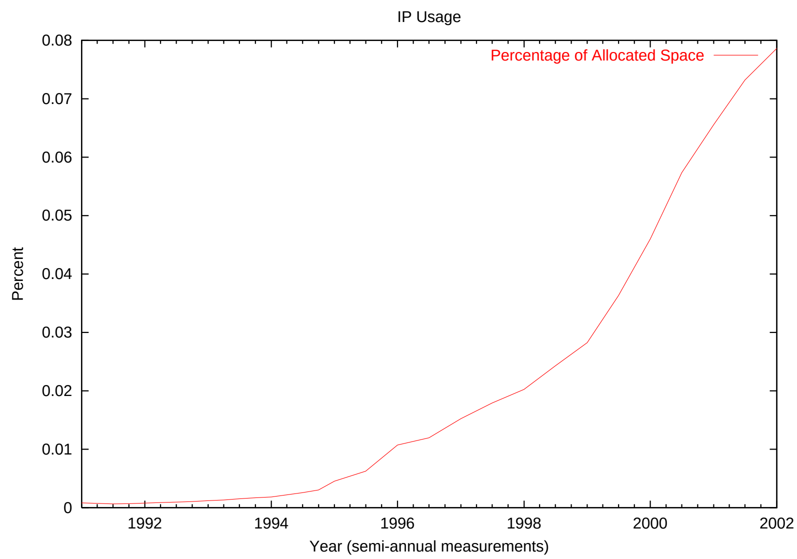

Figure 3.5: Ratio of Hosts (or live IPs) over IP space

3.2.2 Using Host Counts

It should now be apparent that a history of host appearances on the Internet is important. After all, a building's occupancy cannot be determined by its capacity. Recall our two sources of information, Netsizer [30] and IDS [20], of which host counts only from the latter are publicly available.

To better understand how IP addresses are assigned in relation to how they are allocated, we plot the ratio of $\frac{\text{Hosts}}{\text{Allocated IPs}}$ in Figure 3.5. Note that, although this ratio is growing, only 0.08% of all IP space is currently in use. This only reinforces the need for a good host count.

Returning to the desire to understand Internet processes that are necessary to build a relevant model we plot figure 3.6.

Now we can see a more accurate representation of average AS growth. To be clear, this plot shows the appearance of new hosts within an AS over time. Notice

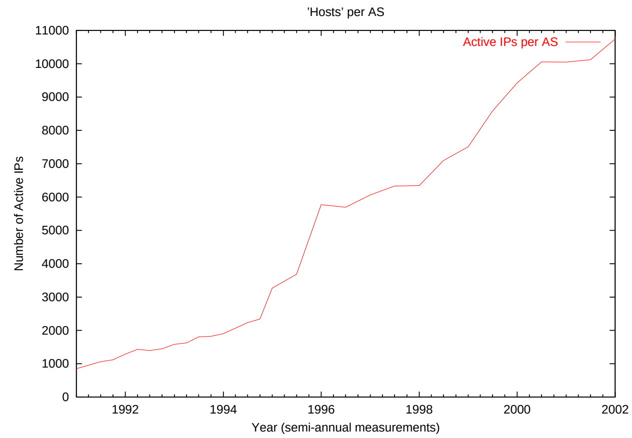

Figure 3.6: Ratio of Hosts (or live IPs) per AS, $\frac{Hosts}{AS}$

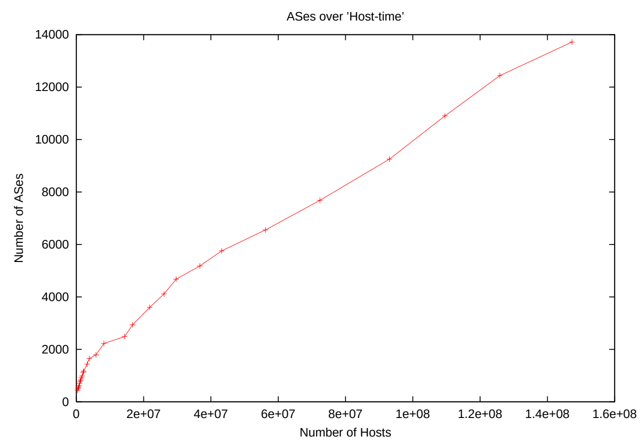

Figure 3.7: Growth of AS wrt Appearance of a Host ("host time")

the line seems near-linear.

But Figure 3.6 is skewed by astronomical notion of time. Hence we plot in Figure 3.7 the number of ASes in *host-time*, where events are appearances of new hosts rather than passing of days and months. The linear relationship hinted at in the previous plot is now only more apparent and gives us reason to proceed further.

Chapter 4

Constructing a Simple Model

We motivate our model with observations on the growth of ASes and hosts over time. Using AS number allocations collected from the three Internet registries and estimates of the number of Internet hosts collected from the Internet Domain Survey [20] , we plot the growth of ASes and Internet hosts over the last decade in Figure 4.1. As one might naturally expect, both plots give evidence of exponential growth.

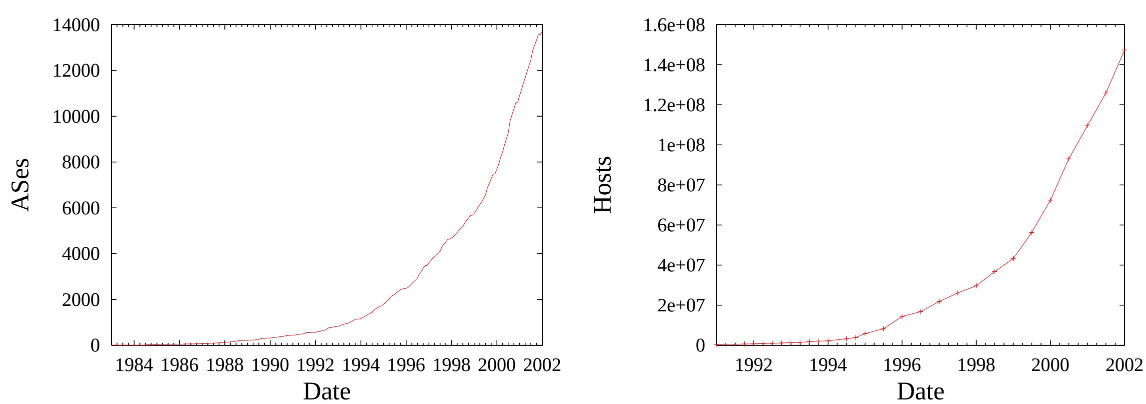

Figure 4.1: Growth in the number of a) ASes and b) Internet hosts.

These observations provide starting points for our model.

4.1 A Simple Model and its Analysis

We summarise the model provided in [18] constructed using the observations made regarding exponential growth.

Let $N(t)$ be the total number of ASes ('N' stands for 'number') and $M(t)$ be the total number of hosts ('M' stands for 'mass') in the system. The simplest growth model consistent with the observations above is mathematically described by linear equations

$$\frac{dN}{dt} = qN, \qquad \frac{dM}{dt} = pM + qN. \tag{4.1}$$

Here q is the rate of creation of new ASes and p is the rate of creation of new nodes. When a new AS is created, the host is given that new label, explaining the qN term in Eq. (4.1). (For now, we shall assume that there is no merging of ASes; moreover, we assume that links do not affect growth processes and hosts and links never disappear.) Solving for N and M gives

$$N(t) = N(0)\, e^{qt}, \tag{4.2}$$

$$M(t) = A\, e^{pt} + BN(t), \tag{4.3}$$

with A, B being simple functions of the initial data $M(0), N(0)$ and the parameters p and q. At the special point $p = q$ the coefficients diverge ($A = B = \infty$), reflecting that the exact solution is actually a linear combination of e^{pt} and $t\, e^{pt}$. Thus the average AS size $\langle s \rangle \equiv M(t)/N(t)$ could in principle exhibit the following asymptotic behaviors:

$$\langle s \rangle \sim \begin{cases} finite & \text{when } p < q, \\ \ln N & \text{when } p = q, \\ N^{(p-q)/q} & \text{when } p > q. \end{cases} \tag{4.4}$$

We show later (Figure 4.3) that the average AS size grows over time (and with

N), thus the inequality $p > q$ must hold.

Let $N_s(t)$ be the number of ASes with s nodes. This size distribution satisfies the rate equation[1]

$$\frac{dN_s}{dt} = p\left[(s-1)N_{s-1} - sN_s\right] + qN\delta_{s,1}. \tag{4.5}$$

We already know $N(t) = N(0)\,e^{qt}$. Solving Eqs. (4.5) recursively and expressing in terms of N rather than t yields

$$N_s = n_s N + \sum_{j=1}^{s} C_{sj} N^{-jp/q}. \tag{4.6}$$

The coefficients C_{sj} depend on initial conditions while n_s are universal. Asymptotically, only the linear term $n_s N$ matters. To determine this dominant contribution, we insert $N_s(t) = n_s N(t)$ into Eq. (4.5). We arrive at the recursion relation

$$\left(s + \frac{q}{p}\right) n_s = (s-1)n_{s-1} \tag{4.7}$$

for $s \geq 2$, while for $s = 1$ we recover $n_1 = q/(q+p)$. A solution to recursion (4.7) reads

$$n_s = \frac{q}{q+p}\,\frac{\Gamma(s)\,\Gamma\left(2+\frac{q}{p}\right)}{\Gamma\left(s+1+\frac{q}{p}\right)}. \tag{4.8}$$

Asymptotically, the ratio of gamma functions simplifies to the power law,

$$n_s \sim C\, s^{-\tau}, \tag{4.9}$$

with $\tau = 1 + q/p$ and $C = \frac{q}{q+p}\,\Gamma\left(2+\frac{q}{p}\right)$.

[1] In the large time limit, the random variables $N_s(t)$ become highly localized around corresponding average values.

4.2 Estimating Growth Rates

In order for us to validate the proposed growth model, we first need to estimate the parameters p and q, the growth rate of the number of hosts and ASes, respectively.

To estimate these rates, we explored a number of alternatives (detailed in Chapter 3) before selecting the methods we deemed most appropriate. For example, BGP tables appear to be a viable alternative for estimating both rates; however, logs only date back to around 1997; moreover, not all IP addresses within a prefix advertised in a BGP prefix are actually in use. The best available method seems to be to use the publicly available routing number allocations provided by the ARIN, RIPE, and APNIC registries. Each keeps a public record dating back to the early 1980s of routing number allocations which include, among other details, the routing number and its type (IP or AS), the date on which the number was allocated, the quantity (and in the case of IP allocations, the starting address). It should be noted that RIPE does not publish AS number allocations, though many allocations to that region have been recorded by ARIN.

From these tables we derived the plot of AS growth in Figure 4.1(a), and plotted again on logscale in Figure 4.2(a). Here we assume that an AS typically comes into existence on the Internet shortly after it is allocated, thus the allocations provide a good estimate for q. Also recall that we are primarily interested in the overall *rate* of growth. Fitting this logscale plot to a line reveals that AS numbers are indeed allocated at an exponentially growing rate. We then estimate q by the slope of the linear regression fit to the curve, or approximately $3.8 \cdot 10^{-4}$.

Estimating p, however, is more difficult. As noted in Chapter 3, allocations of IP addresses are made in bulk by Internet Registries; hence many more IP addresses are allocated than the number of IP addresses in use. The Registries' allocations

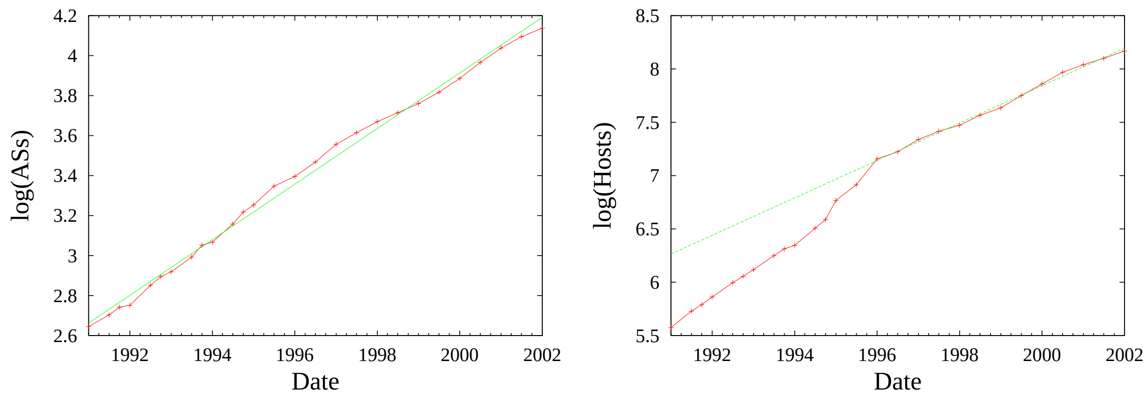

Figure 4.2: Left a) AS Growth, Right b) Host Growth.

statistics show that approximately 50% of IP addresses have been allocated and that the number of allocated IP addresses is growing much less than exponentially, as can be seen in Figure 3.3. We conjecture this trend does not result from the Internet growing less than exponentially, but rather from a growing tendency to better manage allocated IP space.

The Registries provide an excellent record of AS births, but it is infeasible to record IPs in use (and hence, record births of hosts and routers on the Internet). The records of host growth we considered were Telcordia's Netsizer [30] (no longer a publicly available service) and the widely cited Internet Software Consortium's "Internet Domain Survey" project. The host count they develop is based on a reverse DNS process; details can be found at [20]. We can be certain that Registry IP allocation records do not provide host growth statistics since (using IDS numbers) usage of allocated IP space has increased from less than 1% in 1994 to 8% in 2002.

Using the numbers published by IDS, we plot host growth in logarithmic scale in Figure 4.2(b). This plot seems to show a change in slope around 1996. Using the more conservative growth rate, *i.e.* the best fit line of the curve following 1996, we find p to be about $4.8 \cdot 10^{-4}$.

We emphasize that while host count may well underestimate the actual number of hosts on the Internet, we are primarily interested in estimating the slope of the curve; our model is unaffected by scaling factors.

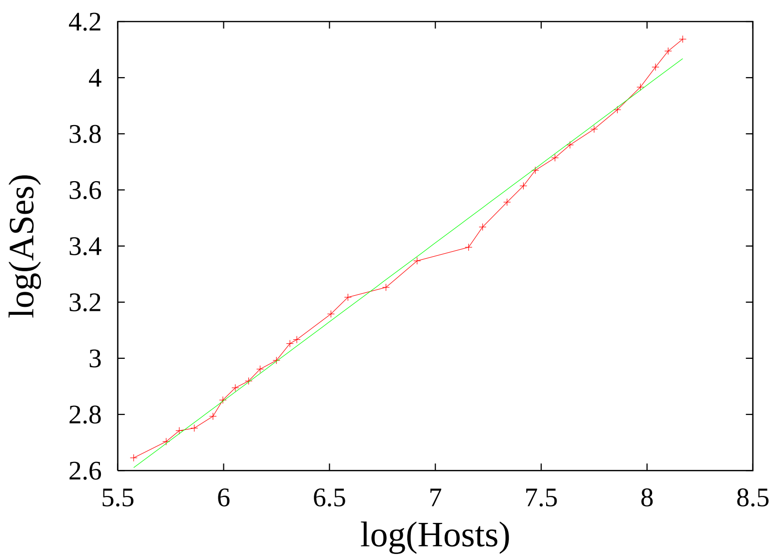

Figure 4.3: Number of Hosts vs. Number of ASes

4.2.1 Analytical Validation

Our model also makes a very specific prediction on the relationship between N and M as the system evolves. From Eqn. 4.4, we expect

$$\langle s \rangle \equiv M(t)/N(t) \sim N(t)^{(p-q)/q},$$

i.e. $M \sim N^{1+(p-q)/q}$. Indeed, we see clear evidence of a power-law fit between M and N when we plot their relationship on log-log scale in Figure 4.3. The predicted slope is 1.26 and the slope of the linear regression is 0.56, so while the model is in the right ballpark, some additional investigation is warranted.

4.2.2 Empirical Validation

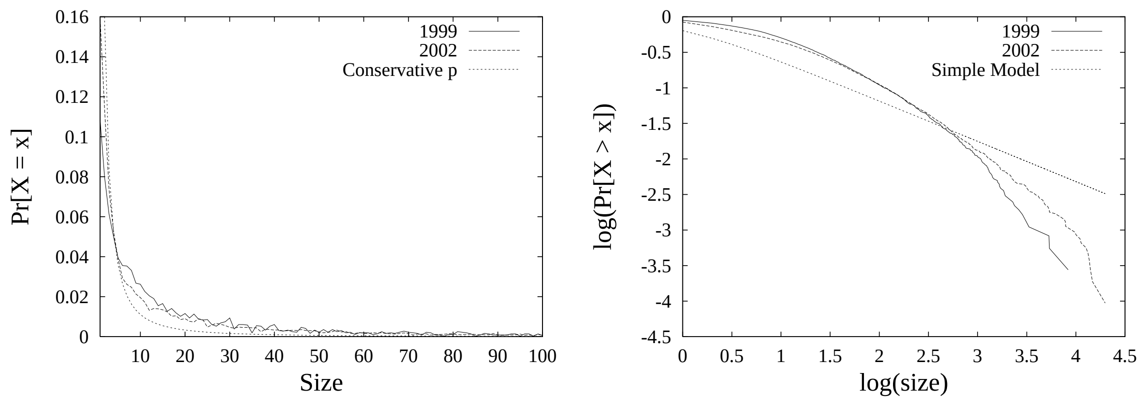

Figure 4.4: Simple Model Predictions vs. Measurements

We first compare the model's predictions using our estimates for p and q against empirical data from 1999 and 2002 in Figure 4.4. Size distributions drawn using Mercator data in 1999, and Skitter data in 2002. The pdf is provided to show rough agreement in the body of the distribution, the log-log plot of the ccdf shows the quality of the fit in the tail. We describe our efforts to improve the simple model's predictive power and accuracy next.

Chapter 5

A More Complete Model

5.1 Modelling AS Mergers

The model described in Section 4.1 is appealing because of its simplicity, but fails to account for a set of prominent events in our datasets — namely, *mergers* between pairs of ASes. In our datasets, we observe these mergers, or coalescence events, in our BGP logs when we witness one AS begin to advertise the set of IP addresses formerly advertised by another AS that then disappears. We provide our methodology for detecting these events in full detail in Section 5.3.1. Coalescence markedly impacts the manner in which ASes grow, since they enable an AS to grow by a multiplicative factor at a single timestep. In this chapter, we describe how to augment the model to incorporate mergers, analyze the asymptotic behavior of the model and its predictions, and compare the predictions to measurements derived from our data sets.

Now, we shall take into account the increase of the total number of nodes, labels, and the merging between different labels, using the same assumptions as before. Recall the notation introduced in Chapter 4. The model is now described by linear equations

$$\frac{dN}{dt} = (q - r)N, \qquad \frac{dM}{dt} = pM + qN.$$

The new parameter, r, is the rate of coalescence, i.e. the rate at which two ASes decide to merge. As before, solving for N and M gives differential equations:

$$N(t) = N_0\, e^{(q-r)t} \quad \text{and} \quad M(t) = A\, e^{pt} + BN(t). \tag{5.1}$$

Following the same analysis of asymptotic behavior as in Section 4.1, and reasoning as before that the average AS size is large and growing, the inequality $p > q - r$ must hold, which implies $\langle s \rangle \sim N^{(p-q+r)/(q-r)}$.

5.2 Implications of the Model

Let $N_s(t)$ be the number of ASes with s nodes. This size distribution satisfies the rate equation

$$\begin{aligned} \frac{dN_s}{dt} &= p\,[(s-1)N_{s-1} - sN_s] + qN\delta_{s,1} \\ &+ \frac{rN}{K}\left\{\sum_{i+j=s} K_{ij}N_iN_j - 2N_s\sum_{j=1}^{\infty} K_{sj}N_j\right\}. \end{aligned} \tag{5.2}$$

The first term on the right-hand side accounts for growth that proceeds with rate p: When a node is added to an AS with $s-1$ nodes, the number of ASes with s nodes increases by one; similarly when a node is added to an AS with s nodes, the number of ASes with s nodes decreases by one. The next term on the right-hand side of Eq. (5.2) accounts for nucleation, with rate q, of new ASes (of size one; one can also study more general situations, e.g., sizes of new ASes can be drawn from a distribution). The last term describes coalescence that proceeds with rate r. This term contains a symmetric "kernel" K_{ij}, the rate of merging between ASes with i and j nodes; $K(t) = \sum_{i,j\geq 1} K_{ij} N_i(t) N_j(t)$ is the proper normalization factor.

Our ongoing work focuses on identifying which kernel most accurately reflects actual coalescence behavior. In what follows, we outline the derivation of the asymptotic behavior of the simplest kernel (an exact analysis is provided in the appendix) and briefly motivate a more general class of kernels, which we can also analyze.

5.2.1 Constant Kernel

Setting $K_{ij} = 1$ transforms Eq. (5.2) into

$$\begin{aligned} \frac{dN_s}{dt} &= p\left[(s-1)N_{s-1} - sN_s\right] + qN\delta_{s,1} \\ &+ \frac{r}{N}\left\{\sum_{i+j=s} N_i N_j - 2NN_s\right\}. \end{aligned} \tag{5.3}$$

Equations (5.3) can be solved recursively. For instance, the number of ASes of the smallest possible size evolves according to $\dot{N}_1 = qN - (p+2r)N_1$. A solution to this equation is a linear combination of two exponents. Asymptotically, the solution simplifies to $N_1(t) = n_1 N(t)$ with $n_1 = q/(p+q+r)$. Similarly, each $N_s(t)$ grows linearly with N. Writing

$$N_s(t) = n_s N(t), \tag{5.4}$$

we recast Eq. (5.3) into the recursion relation

$$\begin{aligned} (q-r)n_s &= p\left[(s-1)n_{s-1} - sn_s\right] + q\delta_{s,1} \\ &+ r\sum_{i+j=s} n_i n_j - 2rn_s. \end{aligned} \tag{5.5}$$

Further discussion and analysis is available in [18].

5.3 Estimating the Rate of Coalescence

Unfortunately, there is no obvious means of tracking AS mergers on the Internet, since we are not aware of any publicly available records providing this information. We therefore resort to making inferences, specifically, by examining aggregated BGP table archives stored by RouteViews [32] at U. Oregon, NLANR and PCH since 1997. Our strategy is to to identify signatures of these merger events from comparisons of sets of daily BGP snapshots This strategy is complicated by the presence of considerable daily churn [9], clouding events of interest with substantial noise.

We argue that aside from churn, there are two reasons for an AS to disappear from daily BGP snapshots.

Coalescence: The IP prefixes formerly advertised by one AS are now advertised by a different AS, and the former AS disappears from BGP tables.

Evaporation: The IP prefixes formerly advertised by one AS simply disappear from the BGP tables.

Note that our methods cannote detect AS mergers in which the acquiring AS retains use of the acquired AS number as well as its own.

To infer these two events, we first identify all "suspicious" events on consecutive daily BGP snapshots. We define a suspicious event to be either 1) the occurrence of identical IP prefixes advertised by two different ASes on successive days, or 2) an IP prefix advertised by one AS on one day, followed by a day in which that exact prefix does not appear, moreover the longest matching prefix including the missing prefix is advertised by a different AS. Of course, many of these suspicious events are due to normal BGP churn and its attendant causes. Therefore, the remaining obstacle is to distinguish an actual merger or disappearance from an instance of BGP churn. It is difficult to distinguish coalescence from evaporation using our methods, since an AS whose IP prefixes evaporate into a larger block of address space is indistinguishable from an AS which coalesces with the AS advertising the surrounding IP block(s).

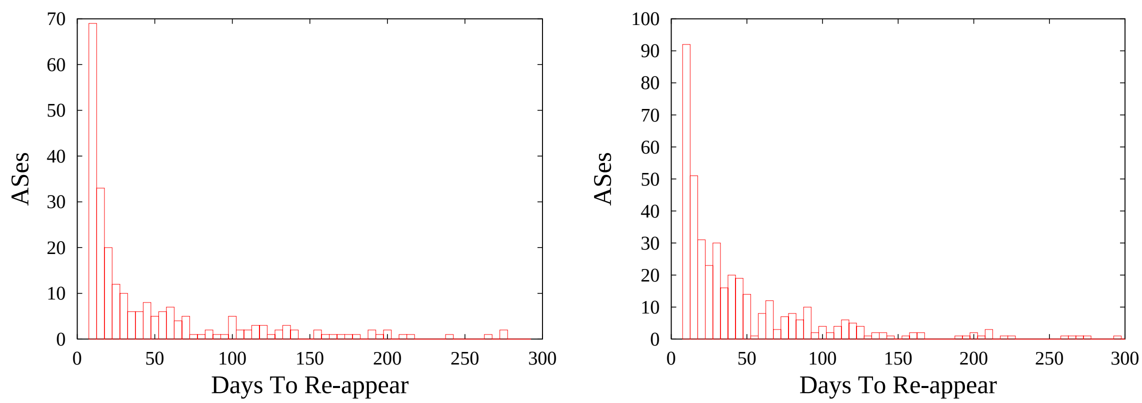

Figure 5.1: AS Reappearance Time. Left: 2000, Right: 2001.

The distinguishing characteristic we use is to consider the *duration* of time that the AS associated with one or more suspicious events actually disappears from the BGP logs (our concern is that AS numbers are eventually reused). To determine an appropriate cut-off threshold, we measured the duration it takes in days for an AS number to reappear after failing to advertise all of its IP blocks, using RouteViews' BGP tables spanning 01/02/2000 - 11/29/2000, and 02/20/2001 - 02/28/2002. Figure 5.1 presents histograms (bin width = 5 days) of the time taken for an AS number to re-appear once it has relinquished its IP space. For clarity, we remove the first bin, which clearly corresponds to BGP churn and constitutes the overwhelming majority of disappearances. In total, 89% and 79% of ASes reappeared in the 2000 and 2001 datasets, respectively, and the majority of these returned within a few days of disappearing. For this reason we feel it is reasonable to assume that beyond a cutoff of between three months to a year, the suspicious event is not due to churn. We record a suspicious event as a merger when,

- there is a handover of IP space as discussed above, and
- the AS number losing IP prefixes then disappears, and
- the AS number does not return in the observed interval.

Applying this analysis to the BGP data allows us to measure the quantities needed to estimate r: $L(t)$ is the total number of ASes present in the tables at time t; and $C(t)$ is the total number of ASes that have merged into another AS by time t. Then, using Eqn. (5.1),

$$\begin{aligned} L(t) \sim e^{(q-r)t} &\sim (L(t)+C(t))/e^{rt} \\ L(t)/(L(t)+C(t)) &\sim e^{-rt} \end{aligned}$$

5.3.1 Validation of the Merger Model

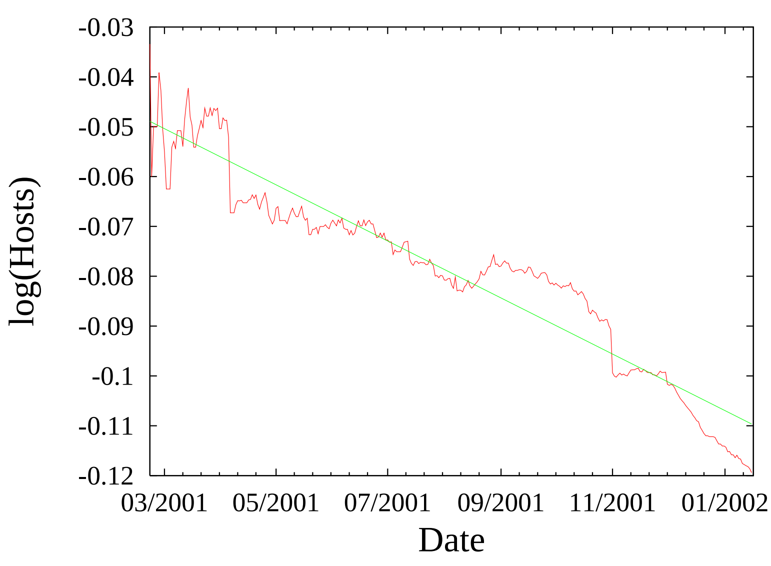

Figure 5.2: Rate at which ASes Merged in the 2001 Dataset.

In Figure 5.2 we plot $L(t)/(L(t)+C(t)$ on semi-log axes, yielding an estimate $r \approx 1.8 \cdot 10^{-4}$. Solving for n_s in Equation 5.5 allows us to summarize the size distribution predicted by the coalescence model in Figure 5.3. As before, only a fit with the body can be seen from the pdf, so a log-log ccdf is provided. Overall, while we find that the merger model is more accurate than the simple model in predicting the distribution of small to medium-sized ASes, it still does not give an accurate prediction of large ASes in the tail of the distribution. More work is necessary to investigate the nature of the kernel in a merger which we hope will allow us to better characterise merger events.

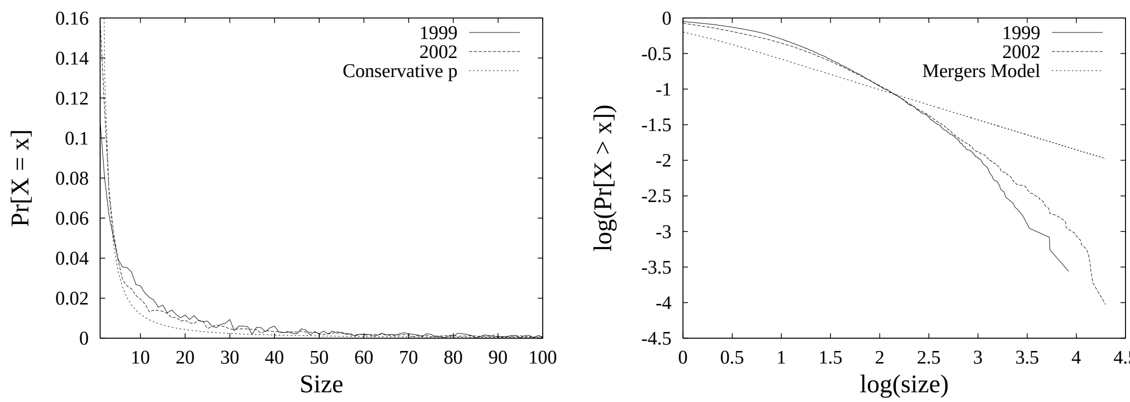

Figure 5.3: Coalescence Model Predictions vs. Measurements

Having discussed in detail two models to describe AS growth processes, we now turn our attention to the influence these processes exhibit upon the structural development of inter-AS structure, specifically with respect to AS degree.

Chapter 6

AS Degree Formation

The previous section showed that a power-law size distribution emerges in the presence of exponential growth of ASes and hosts. In this section we extend that idea to incorporate AS degree.

The key assumption we make is that as an AS grows, it will establish links with other ASes. In this section, we show that if link formation occurs in rough proportion to an AS's growth, then AS degree distribution will show high variability. More precisely, if each time a new node is added to an AS it forms an inter-AS link to some other randomly chosen AS with some fixed probability, then AS degree distribution will show high variability. Furthermore, this need only be in "rough proportion;" for example, we show in Section 7.3 that the result still holds if connection probability varies with the log of the AS size.

Any such link formation process is simple since it only depends on growth, it is flexible since there are no influencing agents other than size, and no global knowledge of other AS degrees is required to make link formation decisions. In fact, no global knowledge of any kind is required.

The simplest interconnection process is detailed in the algorithm below. Recall the notation from Section 4.1 where t is time and $N(t)$ is the number of ASes in the system. Let $M_i(t)$ be the mass of AS i, t_i be the time AS i is inserted into the system, and x be some fixed probability. At each timestep t two kinds of events occur: some new ASes are born, and existing ASes grow. Starting at $t = 1$:

i. Calculate the total number of ASes according to $N(t) = e^{qt}$.

ii. Insert $\lfloor N(t) \rfloor - \lfloor N(t-1) \rfloor$ new ASes with a size of 1 and out-degree of 1, where the neighboring AS is chosen uniformly at random.

iii. Calculate the number of total routers within AS i according to $M_i(t) = e^{p(t-t_i)}$.

iv. For each AS i, insert $\lfloor M_i(t) \rfloor - \lfloor M_i(t-1) \rfloor$ new routers. Each new router creates an inter-AS edge with probability x, and if an edege is created, then invoke a *select* operation to determine to whom the new AS-to-AS link is created.

The *select* operation is left unspecified to emphasize the flexibility of the link formation process and its dependence only on the AS size. We consider only the simplest selection operation, where a target AS is chosen uniformly at random.

Even though this is a random connection process, ASes that are larger in size will also have higher degree. Thus, the degree distribution that results should be highly variable. We show in the following sections that a highly variable degree distribution does result, and that this distribution fits well when compared against distributions observed in the Internet.

Chapter 7

Validation

We validate our analysis and simulation results against empirical degree distributions in the following sections.

7.1 Empirical Data Sources

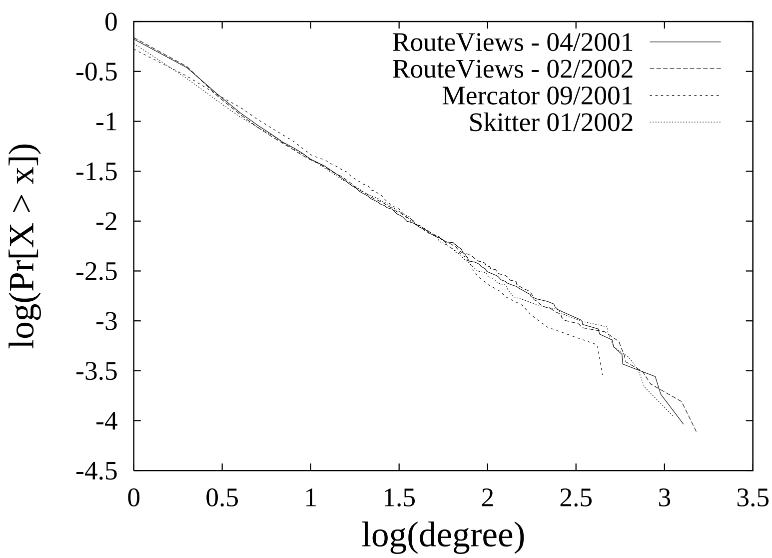

Figure 7.1: Degree Distributions Inferred from 4 Sources.

Before we can discuss our model's influence on inter-AS connectivity, and the validity of the results, we first discuss the empirical observations themselves. There are a number of sources from which we can draw AS-level degree distribution. We infer empirical degree distribution through two distinct methods, applied to three different sources.

7.1.1 BGP Adjacency

The first method is to infer AS degrees from BGP tables. For this purpose we use BGP tables from the RouteViews project [32] collected in April 2001 and February 2002. Recall from Section 3.1.2 that an entry in a BGP table consists of an IP block represented by its prefix, followed by a sequence of ASes (an AS path) that must be traversed to reach an IP address within that range. We can infer an adjacency in the AS-level graph for a pair of ASes whenever they appear in succession within any path.

While this inference method typically avoids false positives (adjacencies which are not actually present, but appear to be present), it suffers from false negatives, since not all AS adjacencies are advertised across BGP [14].

7.1.2 AS Overlays

A second method for determining AS degrees is to annotate a router-level map with each router's associated autonomous system. Nodes in the router-level graph are labeled using IP addresses. In the overlay produced by annotating the router-level graph, each node is further labeled with its associated AS. The approach is detailed in [13]; we summarize the approach here. An IP is associated with an autonomous system by performing a lookup in BGP tables (archived from the same time period in which the router-level map was collected). First, find the longest matching prefix of an IP address within the BGP table; the last entry in the path vector is the number of the AS which owns that IP address. A complete inspection of every edge in the annotated router-level graph reveals an inter-AS edge wherever any pair of nodes are labeled with distinct AS numbers.

This method has numerous advantages over AS maps inferred from BGP tables directly. It provides an AS map at a finer granularity; aggregated ASes are revealed as are multiple links between ASes. However, this method suffers from the following drawback. Any single BGP table is potentially incomplete and can be limited by path hiding from parent ASes (in order to reduce message and table sizes). Sets of BGP tables are used to reduce the magnitude of this problem, with the belief that more BGP tables reveal more information. However, no AS can observe the existence of another AS which is hidden by its parents.

We draw on router-level maps gathered from the Mercator project [19] in August 2001, and another provided by the Skitter project [33] gathered in January 2002.

7.1.3 Summary of Empirical Observations

Statistics, dates, and sources of all datasets drawn from RouteViews, Mercator, and Skitter are summarized in Table 7.1.3.

Source	ASes	Edges	Date	Method
Route Views	10854	47847	04/01	BGP Adjacenies
Route Views	12875	57385	02/02	BGP Adjacenies
Mercator	3478	13590	08/01	AS Overlay
Skitter	9206	38334	01/02	AS Overlay

Table 7.1: Summary of Data Sources

The degree distributions plotted in Figure 7.1 show that all methods and sources yield similar results. For subsequent comparisons, we use the distribution drawn from the autonomous system overlay constructed using the Skitter dataset collected in January 2002 as a baseline for comparison against simulation results.

7.2 Constant Connectivity Models

Section 4.1 showed that the size distribution that results from our model has a power-law tail. However, since the growth model does not directly describe degree, we turn to simulation to determine the influence of size and growth of the simple model on degree.

The simulation is executed using the algorithm in Section 6 using rates $p = 1.1 \times 10^{-3}$ and $q = 8.7 \times 10^{-4}$ estimated in Chapter 4.2.2. The degree distribution predicted by our model is plotted against observed degree distributions in Figure 7.2. We found empirically that using fixed probability $x = 0.10$ results in vertices of our simulated graphs having a roughly commensurate average degree to that of the Skitter dataset. Where the discrepancy does occur, the general trend is a tendency for our model to underestimate the degree of small to medium sized ASes, while overestimating the degree of larger ASes.

Figure 7.2 shows that the predicted degree distribution is similar to that of the Skitter dataset. Discrepancies can potentially be removed by refining the decision processes used to form AS to AS connections in the model. In the following section, we explore a refined model which accounts for the size of the AS when determining the relationship between growth and link formation.

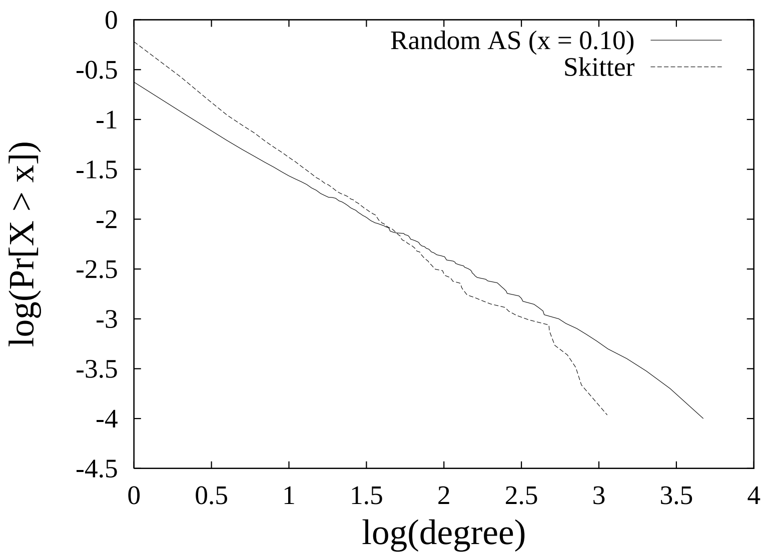

Figure 7.2: Predicted Degree Distribution where $x = 0.10$

7.3 Size-Based Connectivity Models

The previous section shows that link formation occuring with constant probability during growth, while reasonable, could be more accurate. The relationship between predicited and empirical distributions shown in Figure 7.2 suggest that there is room for other practical influences on inter-AS link formation. Here we demonstrate the flexibility of our model by discussing an approach that takes into account the actual size of the AS when choosing to create new links.

We presuppose the following notion: as an AS grows, the ratio between its degree and its size will shrink, and so a constant probability when deciding to create new links may not best relate degree to size. Intuitively, the ratio between the degree of an AS and its size is analogous to the notion of a surface-to-volume ratio. In graph-theoretic terms, this ratio is often referred to as the conductance of a subgraph.

Definition 1 *The* conductance *of an autonomous system i with size (mass) M_i and out-degree d_i is $\frac{d_i}{M_i}$.*

Observations of conductance are estimated from Mercator and Skitter datasets discussed in Section 7.1, and shown in Table 7.3. This table shows that as an autonomous system grows, the average conductance shrinks. While the actual conductance of ASes of a given size varies considerably, this trend holds on average. Note that ASes of size 1 are excluded from the smallest range since an AS of size 1 *must* have conductance of at least 1, and so may bias observations. Also, average conductance in the largest ASes appear to break this trend. This may be an artifact of noise from a small number of data points.

	Data Points		Average Conductance	
Size Range	Mercator	Skitter	Mercator	Skitter
$2 - 10$	1404	4254	0.492	0.866
$11 - 100$	1429	3502	0.242	0.596
$101 - 1000$	359	1050	0.134	0.313
$1001 - 10000$	38	131	0.108	0.213
$10001 - 100000$	1	10	0.20	0.249

Table 7.2: Conductance of ASes

We believe that this decrease in conductance is natural, driven by the decreasing necessity to add inter-AS links as an AS grows. For example, as previously mentioned, an AS of size 1 must have a minimum degree of 1 (otherwise it is not connected to other ASes, and hence cannot be a part of the AS-level map). We speculate that it is more often the case that routers are added to a closed network to increase the capacity and range of the network itself, rather than to connect to other ASes, and so a connection probability that decreases as an AS grows is reasonable.

The ratios and ranges in Table 7.3 show diminishing conductance as AS size increases. To better fit the data observed in Table 7.3, we applied a logarithmic correction factor to implement a "diminishing probability" function, L. This function takes the size of the autonomous system M_i, and a fixed probability x as parameters, and returns a probability value:

$$L(x, M_i) = \begin{cases} x, & \text{when } M_i < 10, \\ \frac{x}{\log_{10}(M_i)}, & \text{otherwise.} \end{cases} \tag{7.1}$$

As before, we use the simple *select* operation which returns a neighboring AS chosen uniformly at random. One point of interest is that L governs only those ASes where size ≥ 10, otherwise the probability of connecting to another AS is artificially inflated for the smallest ASes.

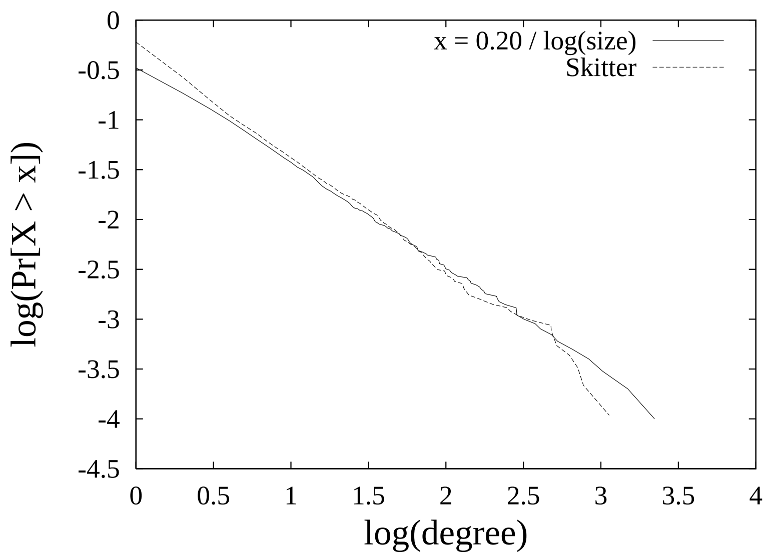

Figure 7.3: Diminishing Probability where $x = 0.20$.

The distribution that results when applying the diminishing probability function is plotted against Skitter data in Figure 7.3, using $x = 0.20$, the value providing the best fit. The two curves are nearly identical, sharing a similar slope, and are virtually indistinguishable throughout the entire body of the distribution.

Chapter 8

Conclusion

Understanding the dynamics of AS size distribution is important both for modeling purposes and understanding connectivity. In this thesis we have proposed two models for the evolution of the AS size distribution. We further explored a model for how highly variable degree distributions may arise in the AS graph as a result of the evolution of AS size.

First, we have provided and analyzed a growth model with rate equations. We have discussed methods for estimating the parameters of this model and shown the size distribution of ASes that it predicts. The model's predictions exhibit size distributions that are in general agreement with empirical data, both in the body and the tail of the distribution. However, discrepancies exist between model and data, particularly in the shape of the tail.

Second, we have suggested that it is important to incorporate the merging of ASes in our models. We show how to do so, and specify the resulting rate equation. The details of this model are highly dependent on assumptions about the manner in which ASes merge, which is captured in the merging kernel – the likelihood that two ASes of given sizes will merge at any timestep. We solve this model for the constant kernel, and show how to estimate the associated parameters. The results point to the need for further analysis of the processes by which ASes merge.

Lastly, we extended the growth model to connect AS nodes. We believe this is the first such size based Internet model for AS-labelled graphs. It is instructive to compare this model with the B-A model.

Like the B-A model, we assume that high variability has arisen via a "rich get richer" phenomenon resulting from an exponential growth process. However the B-A model assumes *preferential connectivity,* meaning that new nodes probabilistically prefer to connect to well-connected existing nodes. Besides requiring that each AS be aware of the degree of each other AS (a strong assumption of global knowledge), the B-A model strongly constrains the resulting connection pattern. This is restrictive; as discussed in [28], many graph realizations are consistent with a given degree sequence, and different realizations may have very different properties. In fact, [27] shows that the AS graph exhibits a high degree of clustering, an effect that is not captured by the particular connection pattern created by the B-A model.

In contrast, the assumption in our model is that *AS sizes* are the underlying cause of high variability, and that a large AS will naturally tend to have a large degree. From this standpoint, our model allows for a much wider range of connection patterns than the B-A model, since the degree of an AS grows as a function of its size, but the choice of *which* AS to connect to can be specified independently, as a separate selection operation. In this thesis we have explored the selection operation in which growing ASes choose peering partners uniformly at random; however we expect that any choice of peering partners that is made without regard to degree (and including those that exhibit a high degree of clustering) will likely show characteristic high variability.

Our results demonstrate that a simple and natural model incorporating exponential growth alone is sufficient to drive both a highly variable AS size distribution and a highly variable AS degree distribution. We motivated this model with datasets that demonstrate exponential growth both in the number of hosts and the number of ASes, and validated the model by comparing the degree distribution our model predicts against observed degree distributions drawn from BGP tables and AS overlay maps. We also provide an analysis of the power-law tail of the AS size distribution that results when our methods are employed.

We have integrated this model into the publicly available BRITE [8, 25] topology generation framework. In future work, one might investigate selection operations that incorporate real-world considerations such as locality, clustering and performance optimization, to provide an even more realistic AS growth model. As part of this effort, better mining techniques of AS time-series data extracted from BGP logs are necessary to better understand the underlying nature of AS growth, interconnection and merging over time [18]. We hope this model is the first in a line of more flexible size-based approaches to Internet modelling.

Bibliography

[1] W. Aiello, F. Chung, and L. Lu. Random evolution in massive graphs. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, Las Vegas, NV, October 2001.

[2] R. Albert and A. Barabási. Topology of evolving networks: Local events and universality. *Physical Review Letters*, 85:5234–5237, 2000.

[3] R. Albert and A. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, January 2002.

[4] D. Alderson, J. Doyle, R. Govindan, and W. Willinger. Toward an Optimization-Driven Framework for Designing and Generating Realistic Internet Topologies. In *ACM HotNets-I*, Princeton, NJ, October 2002.

[5] APNIC. At *http://www.apnic.net/*.

[6] ARIN. At *http://www.arin.net/*.

[7] A.-L. Barabási and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286:509–512, October 1999.

[8] Boston University Representative Internet Topology Generator (BRITE). At *http://www.cs.bu.edu/brite*.

[9] A. Broido, k. claffy, and E. Nemeth. Internet Expansion, Refinement and Churn. *European Transactions on Telecommunications*, To Appear, 2002.

[10] T. Bu and D. Towsley. On Distinguishing between Internet Power Law Topology Generators. In *Proceedings of IEEE INFOCOM*, New York, NY, June 2002.

[11] K. Calvert, M. Doar, and E. Zegura. Modeling Internet Topology. *IEEE Communications Magazine*, 35:160–163, December 1997.

[12] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Towards Capturing Representative AS-Level Internet Topologies. Technical Report UM-CSE-TR-454-02, University of Michigan Computer Science, 2002.

[13] H. Chang, S. Jamin, and W. Willinger. Inferring AS-level Internet Topology from Router-Level Path Traces. In *Proceedings of SPIE ITCom 2001*, Denver, CO, August 2001.

[14] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The Origin of Power Laws in Internet Topologies Revisited. In *Proceedings of IEEE INFOCOM*, New York, NY, June 2002.

[15] A. Fabrikant, E. Koutsoupias, and C.H. Papadimitriou. Heuristically Optimized Trade-offs: A New Paradigm for Power Laws in the Internet. In *International Colloqium on Automata, Languages, and Programming (ICALP)*, Málaga, Spain, July 2002.

[16] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *Proceedings of ACM SIGCOMM*, pages 251–62, Cambridge, MA, September 1999.

[17] M. Fayed, P. Krapivsky, J. Byers, M. Crovella, D. Finkel, and S. Redner. On the Emergence of Highly Variable Distributions in the Autonomous System Topology. *To appear in ACM SIGCOMM Computer Communication Review (CCR)*, 33, July 2003.

[18] M. Fayed, P. Krapivsky, J. Byers, M. Crovella, D. Finkel, and S. Redner. On the size distribution of autonomous systems. Technical Report BUCS-TR-2003-001, Boston University, January 2003.

[19] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery. In *Proceedings of IEEE INFOCOM*, March 2000.

[20] Internet Domain Survey. At *http://http://www.isc.org/ids/*.

[21] P. L. Krapivsky and S. Redner. Organization of Growing Random Networks. *Physical Review E*, 63:066123, 2001.

[22] P. L. Krapivsky, S. Redner, and F. Leyvraz. Connectivity of Growing Random Networks. *Physical Review Letters*, 85:4629–4632, 2000.

[23] D. Magoni and J. Pansiot. Analysis of the Autonomous System Network Topology. *ACM SIGCOMM Computer Communication Review (CCR)*, 31(3):26–37, July 2001.

[24] P. Meakin. Droplet Deposition Growth and Coalescence. *Reports on Progress in Physics*, 55:157–240, 1992.

[25] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An Approach to Universal Topology Generation. In *IEEE MASCOTS*, Cincinnati, OH, August 2001.

[26] A. Medina, I. Matta, and J. Byers. On the Origin of Power Laws in Internet Topologies. *ACM SIGCOMM Computer Communication Review (CCR)*, 30(2):18–28, April 2000.

[27] M. Mihail, C. Gkantsidis, and E. Zegura. Spectral analysis of Internet topologies. In *Proceedings of IEEE INFOCOM*, April 2003.

[28] M. Mihail and N. Visnoi. On generating graphs with prescribed degree sequences for complex network modeling applications. In *Proceedings of Approx. and Randomized Algorithms for Communication Networks (ARACNE)*, 2002.

[29] S. Mossa, M. Barthélémy, H. Stanley, and L.A. Nunes Amaral. Truncation of Power Law Behavior in "Scale-Free" Network Models due to Information Filtering. *Physical Review Letters*, 88(13):138701, March 2002.

[30] Telcordia's NetSizer. At *http://www.netsizer.com/*.

[31] RIPE. At *http://www.ripe.net/*.

[32] RouteViews Project at University of Oregon. At *http://archive.routeviews.org/*.

[33] The Skitter Project. At *http://www.caida.org/tools/measurement/skitter/*.

[34] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, and S. Shenker. Does AS Size Determine Degree in AS Topology? *ACM SIGCOMM Computer Communication Review (CCR)*, 31(5):7, October 2001.

[35] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network Topology Generators: Degree-Based vs. Structural. In *Proceedings of ACM SIGCOMM*, Pittsburgh, PA, August 2002.