The State of Internet Topology Research

Ramesh Govindan

ramesh@icsi.berkeley.edu

Joint work with Hongsuda Tangmunarunkit, John Doyle, Sugih Jamin, Scott Shenker, Walter Willinger

Of Blind Men and Elephants . . .

And so these men of Indostan
Disputed loud and long,
Each in his own opinion
Exceeding stiff and strong.
Though each was partly in the right,
They were all in the wrong!

John Godfrey Saxe

Outline

- What does the Internet topology look like? Discovery
- Can we model it? Modeling
- How does it impact routing and protocols? Impact

Will attempt to summarize evolution of thought in the area, but will not attempt to be objective

Caveats

- Two distinct uses of the term *Internet topology*
  - AS-level BGP peering graph
  - Router-level adjacency graph
- Most analyses consider these as unweighted graphs
  - Not annotated with capacity or latency, for example

Impact

What use is all this?

- Form of neighborhood function qualitatively affects multicast tree scaling [PhillipsShenker99]
- Policy significantly distorts path hop-counts [Tangmuna01]
- Topology can impact protocol performance by up to a factor of two [Radoslavov99]

But beyond that, work in this area just driven by curiosity

Discovery: AS Topology

- From BGP routing information
  - From dumps (most people)
  - From updates [Govindan97], [Alaettinoglu01]
- Augment using Looking Glasses, registries [Chang01]
- From router-level maps [Tangmuna01], [Broido01]

Discovery: Router-level Topology

- Traceroute [Keshav97]
- Trick to resolve router aliases [Pansiot99]
- Web servers to cover address space [Skitter00]
- BGP table to cover address space [Cheswick00]
- Tricks to increase fidelity [Govindan00]
  - Source routing
  - Alias resolution
  - Informed address probing

Discovery: Where are we?

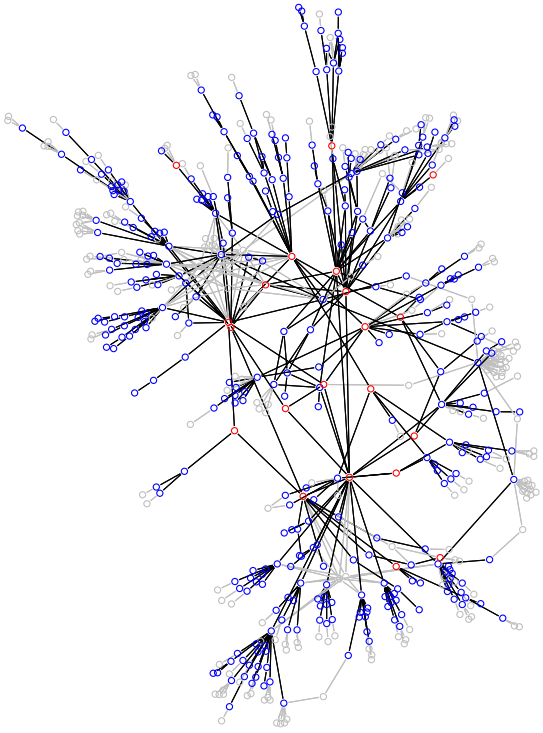

Router-level topology

- Techniques exhibit reasonable fidelity (90%)
- But extent of completeness uncertain
- Collect topology from different points

AS Topology

- About as close as we can get?

Modeling Topology

What is a good topology generator?

Early router-level generators: intuition was physical

- Distance weighted random graphs [Waxman88]
- Structural generators (Transit-Stub and Tiers) with some notion of hierarchy [Doar96] [Zegura97]

Power Laws

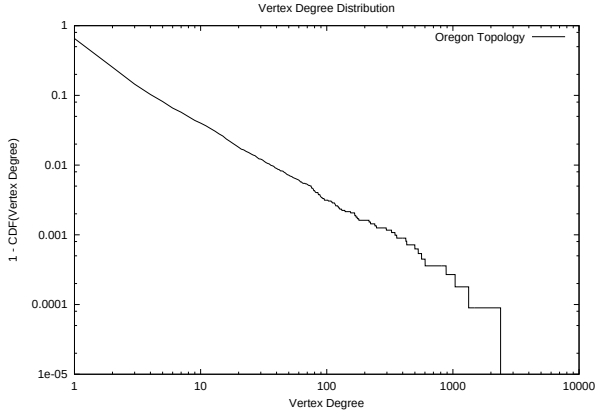

- AS degree power law [Faloutsos99]
- Follow-on results
  - Not quite a power law [Chang01]
  - Actually Weibullian [Broido01]
- Safe to say that the degree distribution is *highly-variable*
- Significant degree variability in router-level graph as well

Modeling Power-Law Topologies

More recent generators based on power-laws

- Empirical models: Inet [Jin00], PLRG [Chung00]
- *Growth* models: BRITE [Medina00]

But, we know the Internet has structure and hierarchy

- . . . How can these degree based generators be right?

Do the growth models really explain AS structure?

Are Degree-Based Generators Right?

What's the right metric?

- Average degree, diameter, eccentricity, average path length [Zegura94]
- Clustering coefficient [Watts98], [Bu01]
- Laplacian spectrum [Vukadinovic01]

What's the methodology?

- Compare metric values for similarly sized graphs (almost everybody)

Our Approach

Metrics

- Large-scale properties are important
- Degree distribution and clustering measure *local* properties
- We don’t know the right metrics, so try them all

Methodology

- Study the scaling behavior using *ball-growing*
- Make *qualitative* distinctions

Caveat: *ad-hoc* and subjective

Key Result

Large-scale properties of Internet graphs well modeled by a power-law random graph (PLRG)! [Tangmuna00]

Topology	Expansion	Resilience	Distortion
Mesh	L	H	H
Random	H	H	H
Tree	H	L	L
Complete	H	H	L
Linear	L	L	L
AS, Router-Level, PLRG	H	H	L
Tiers	L	H	L
Transit-Stub	H	L	L
Waxman	H	H	H

But what about hierarchy?

Hierarchy Metric

In the AS graph, there are well-defined hierarchical relationships (*e.g.*, provider-customer)

- Some work on *inferring* these [Gao00] [Lakshmi01]

In a general graph, what do we mean by hierarchy?

- Some links are more *important* that others
- Paths go *up* the hierarchy, then *down*

We define a link importance metric

- How many nodes would be “upset” if the link was cut
- Look at distribution, as well as the up-down fraction of paths

Hierarchy Findings

Power-law random graphs, AS and router-level graphs have some hierarchy

- But not as strict as a tree

Hierarchy in power-law graphs arises from the degree distribution!

How Does Power Law Arise?

- One hypothesis: an incremental growth model, with preferential connectivity [BarabasiAlbert00]
- This gives rise to a power-law degree distribution
  - Can possibly be tweaked to yield highly-variable distributions
- Model fits nicely, but is it the explanation?

Cause or Effect?

- In practice, is an AS more likely to connect to another AS because it has high degree . . .
- Or because the other AS
  - Has a greater topological reach
  - Has a better engineered backbone
  - Offers deep discounts
  - Has a good reputation for availability
- In other words, is degree the cause or the effect ?

Size Distributions

- A clue: other occurrences of highly variable distributions
  - Companies by size of income or assets
  - Cities by population size
  - Oil reserves by size
  - File or Web document sizes
- *Highly-variable size distributions are ubiquitous!*

Occam's Razor

So, is the highly-variable degree distribution simply determined by a highly-variable AS size distribution?

- UUNet has a large degree because it is a large corporation, not vice versa
- Minnesota Regional Network has a small degree because it has a small network and doesn't have global reach, not vice versa

Caveat: We don't have explanations for skewed size distributions

- . . . so, if true, we have exchanged one mystery for another!

Findings

- AS sizes are highly-variable
- Size and degree are well-correlated
- Highly-variable size distributions are ubiquitous

Size determines degree in AS Topology? [Tangmuna01b]

Wrapping Up

What we know

- AS and router-level topologies have a highly-variable degree distribution
- A power-law random graph models the large-scale structure of these networks
- For the AS graph, size determines degree?

What we don't know

- How does router-level graph arise? (Global optimization?)

Other Topology Work

http://topology.eecs.umich.edu/