The Internet Is Like A Jellyfish

Michalis Faloutsos

UC Riverside

Joint work with:

Leslie Tauro, Georgos Siganos (UCR)

Chris Palmer(CMU)

Big Picture: Modeling the Internet

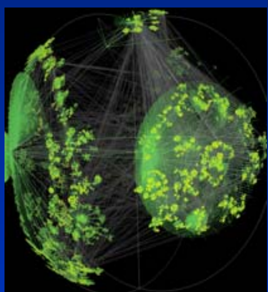

Topology

Traffic

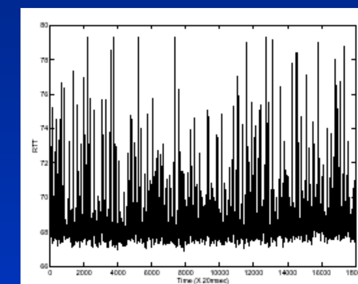

Protocols

Routing, Congestion Control

- **Measure and model each component**
  - Identify simple properties and patterns
- **Model and simulate their interactions**

The Goal of Internet Modeling

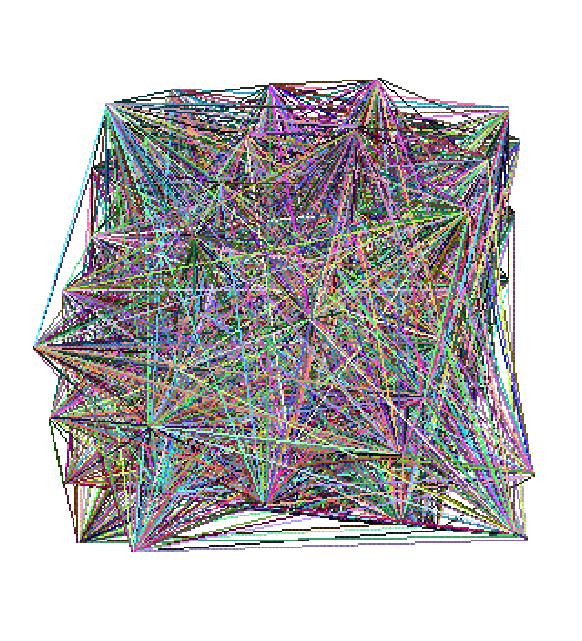

A real Internet instance

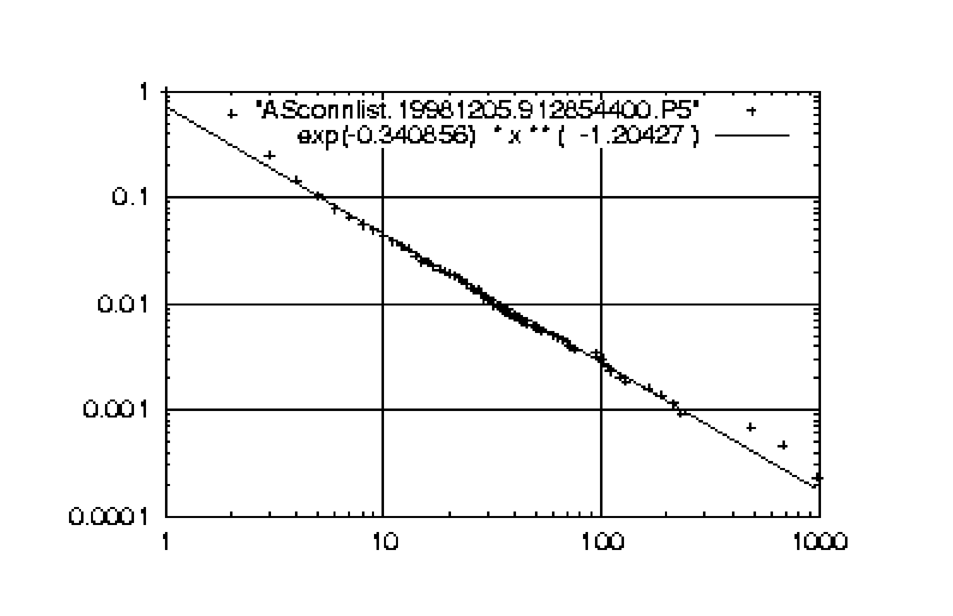

Power-law: Frequency of degree vs. degree

- **Find simple fundamental properties**
- **Understand why they appear and their effects**

Claim: We Need The Right Tools

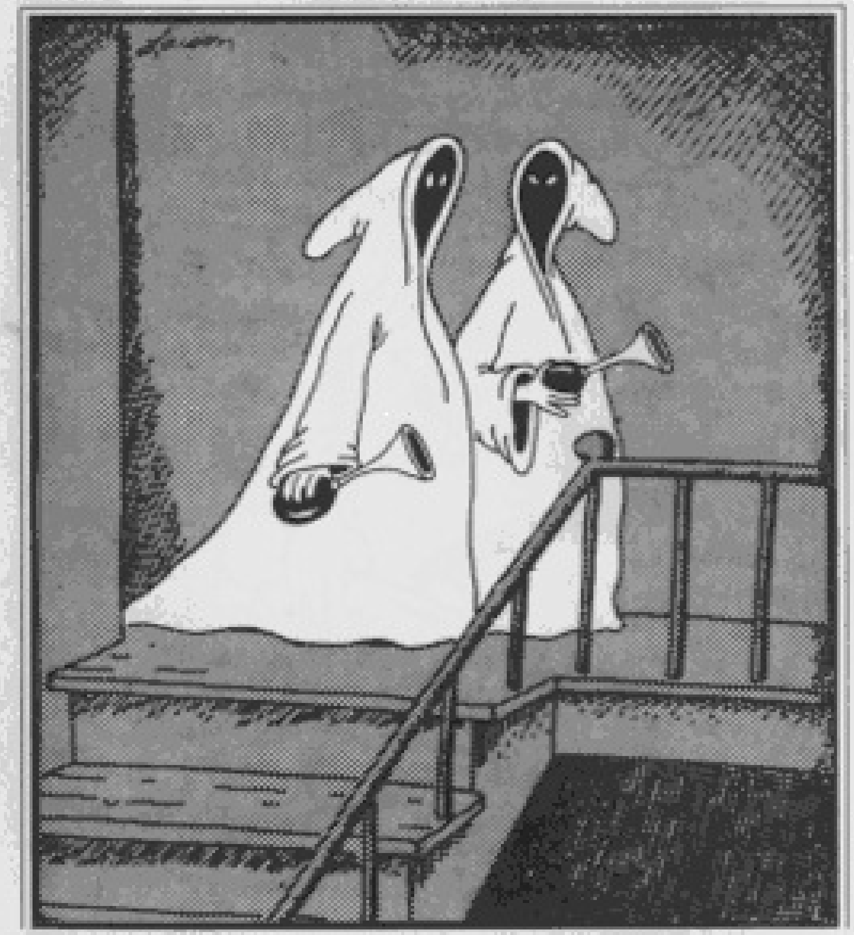

"This is just not effective . . . We need to get some chains."

"This is just not effective…

We need to get some chains"

The Far Side -- G. Larson

The World Wide Web is a Bow-Tie

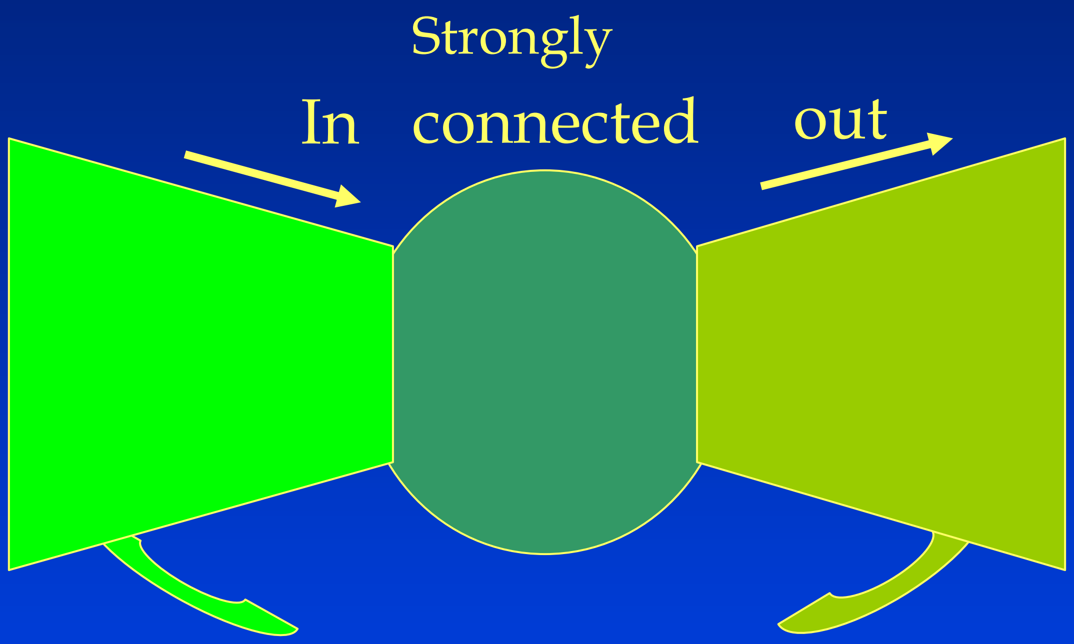

- Core
- In
- Out
- Tendrils

- **Captures several properties [WWW-Tomkins et al]**
- **The components are of comparable size**

The Accuracy-Intuition Space Of Models

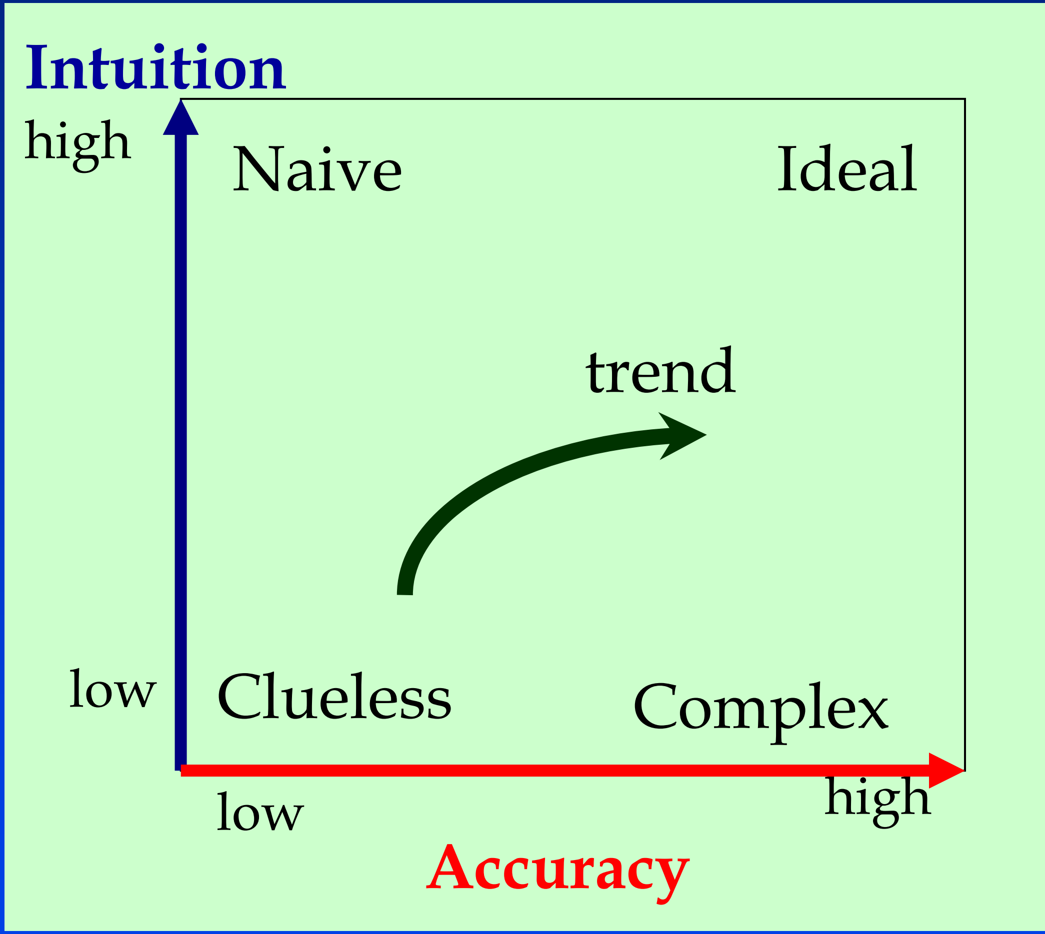

- **More tools…**
  - Self-similarity
  - Power-laws
  - Wavelets
  - Eigenvalues
- **…less intuition**
  - Something a human can picture
- **Is it a real conflict?**

Why Do We Need an Intuitive Model?

- **Human mind is simple**
- **Visualizable: creates a mental picture**
- **Memorable: captures the main properties**
- **Maximizes *information/effort* ratio**
- **Makes you think**

What does the Internet look like?

- **Can I develop a simple model of the AS Internet topology that I can** draw by hand**?**
- **Can I identify a sense of hierarchy in the network?**

Focus: Autonomous Systems topology
and data from NLANR

Possible Topological Models

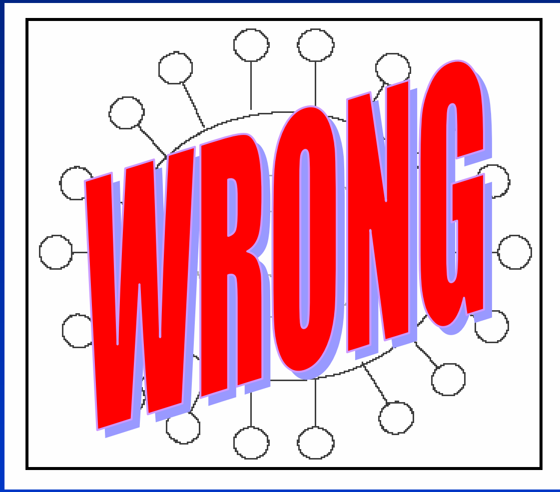

Furball

(One-degree nodes are at the periphery)

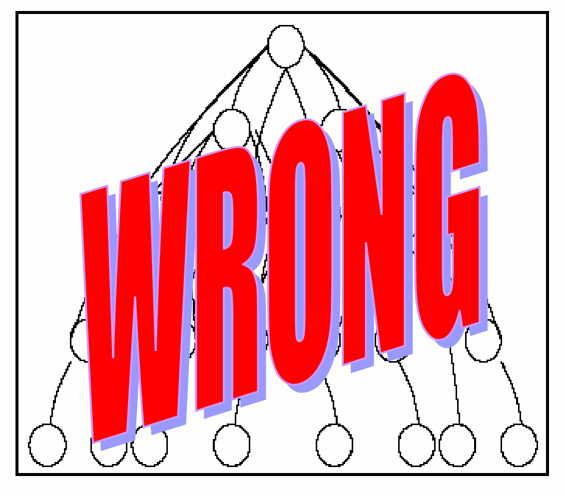

Broom

(Military Hierarchy)

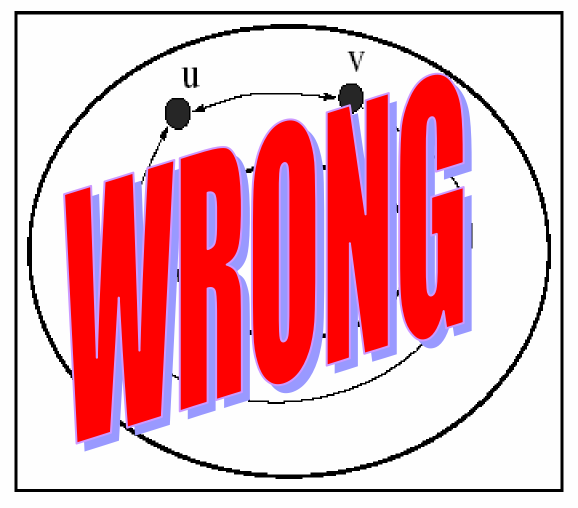

Donut

(Circular connectivity: Around the earth?)

An Intuitive Model : The Internet Topology as Jellyfish

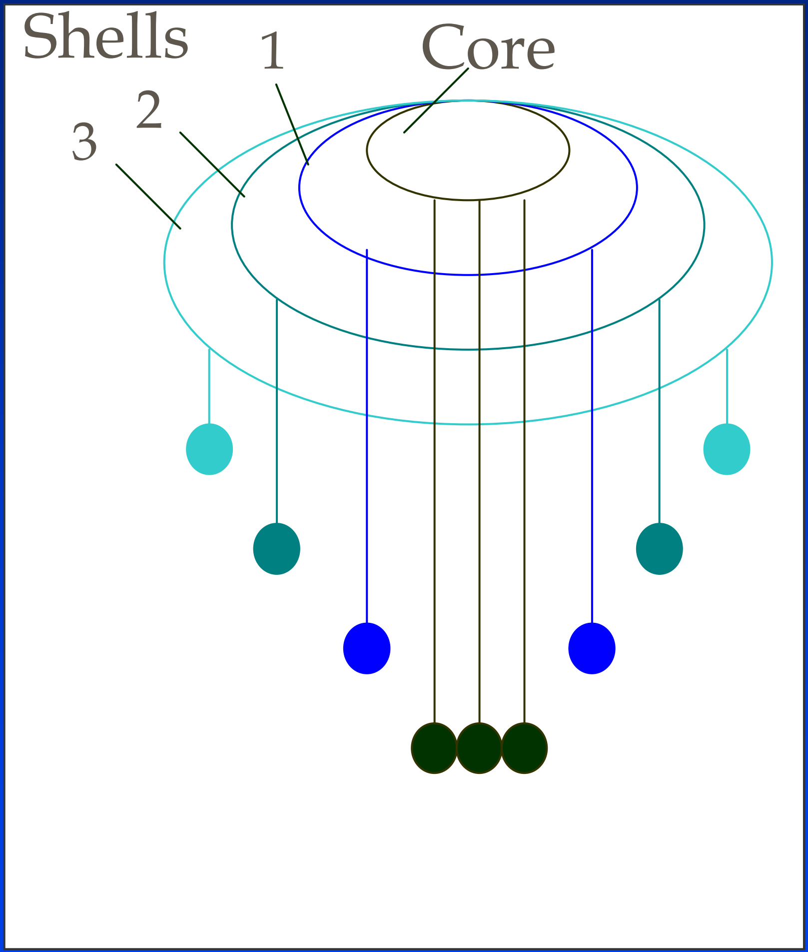

- **Highly connected nodes form the core**
- **Each Shell: adjacent nodes of previous shell, except 1-degree nodes**
- <u>Importance</u> **decreases as we move away from core**
- **1-degree nodes hanging**
- **The denser the 1-degree node population the longer the stem**

Roadmap

- Identify a Hierarchy
  - Define the Importance of a node
- Present topological properties
- Present the jellyfish model
- Why is the jellyfish a good model?
- Conclusions
- Appendix: Latest News on power-laws for the Internet topology

How Can We Develop a Simple Model?

- **We need an anchor and a compass**
- **Anchor:**
  - We need a starting point in the network
- **Compass:**
  - We want to classify nodes according to **importance**

Defining the Importance of a Node

- The topological importance has many aspects
- **Degree:** number of adjacent nodes
- **Eccentricity:** the maximum distance to any other node

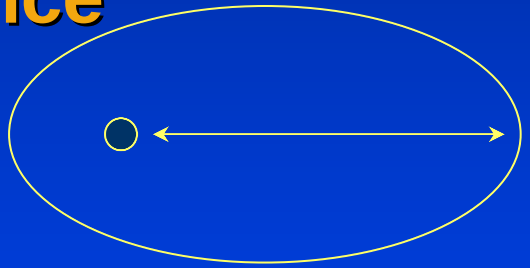

- **Significance:** Significant nodes are near :
  1. many nodes
  2. significant nodes

Significance of a Node

- **The significance of a node is the sum of the significance of its neighbors.**
- **The iterative procedure converges**
  - At each round, total significance is normalized to 1
- **This is equivalent to [Kleinberg '96]:**
  - the eigenvector of the max eigenvalue of the adjacency matrix
- **Relative Significance: Signif. times No. Nodes**
  - Relative Significance = 1, fair share of significance

Roadmap

- **Identify a Hierarchy**
  - Defining the Importance of a node
- **Present topological properties**
- **Present the jellyfish model**
- **Why is the jellyfish a good model?**
- **Conclusions**

Observation 1: Significance and Eccentricity Are Correlated

Significance

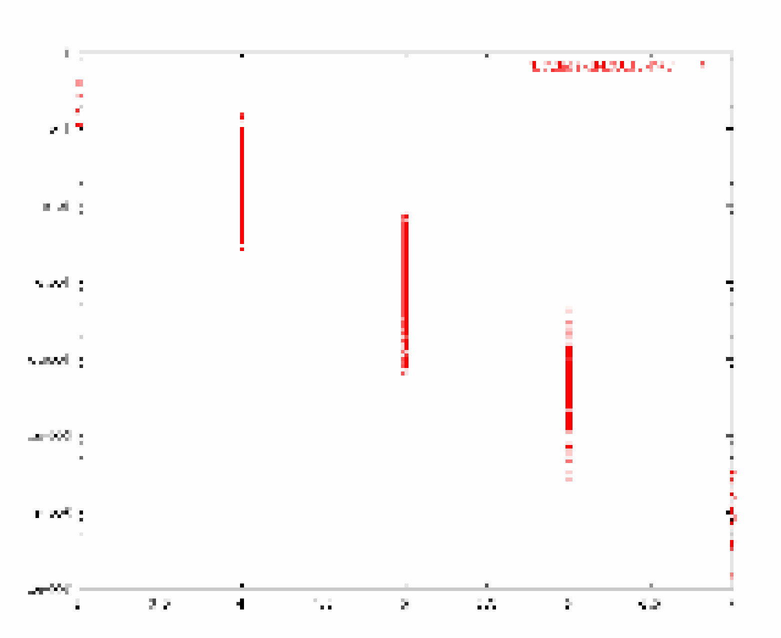

Eccentricity

- Significant nodes have low eccentricity
- Intuitively, significant nodes are in the middle of the network [Global Internet '01]

Observation 2: Many One-Degree Nodes Connect to High-Degree Nodes

Number of 1-degree neighbors

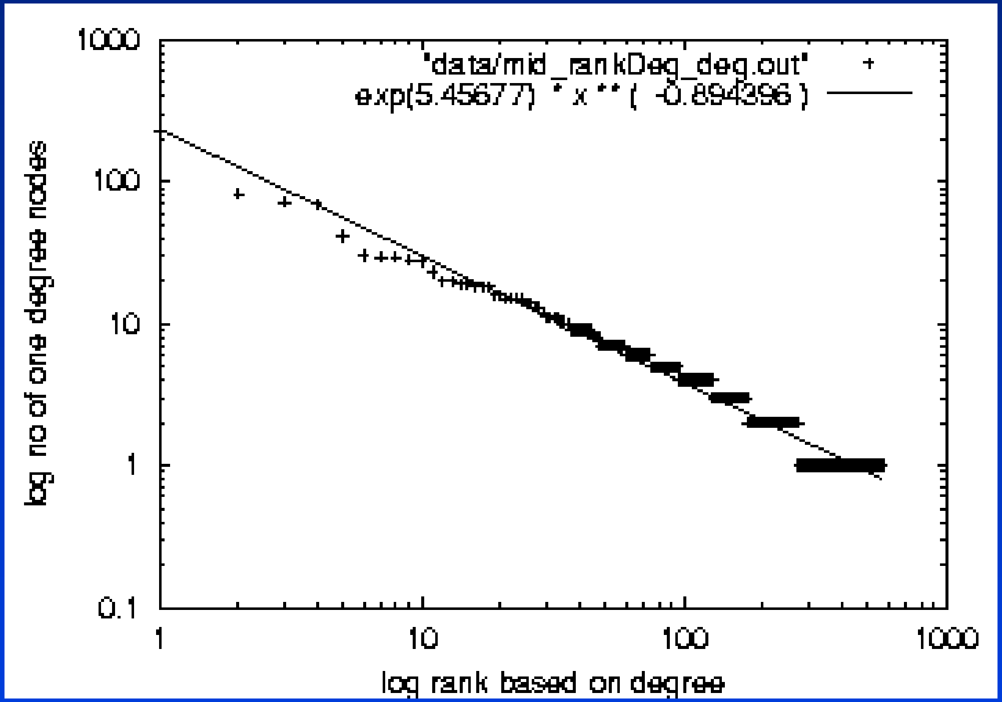

The failure of Furball model

Order of decreasing neighbors

- One-degree nodes are scattered everywhere
- The distribution of one-degree nodes follows a powerlaw

Observation 3: The Internet Premise: One Robust Connected Network

Size of Largest Connected Component

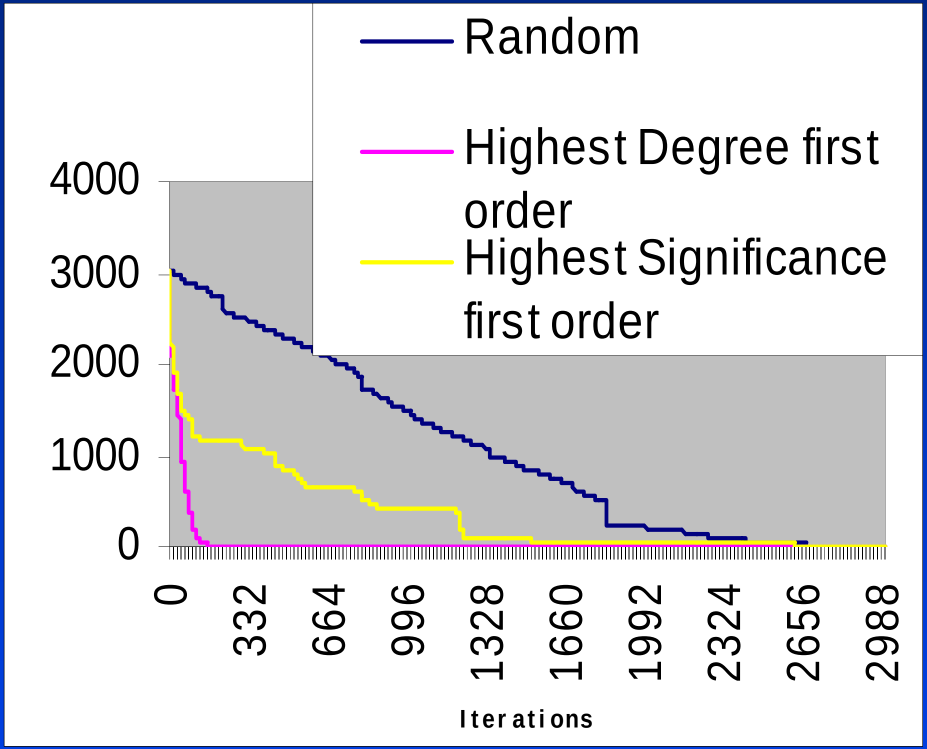

#Deleted nodes

- Robust to random, sensitive to focused failures
- The network tends to stay as one connected component

Observation 4: The Number of Alternate Paths Between Two Nodes

Number of paths

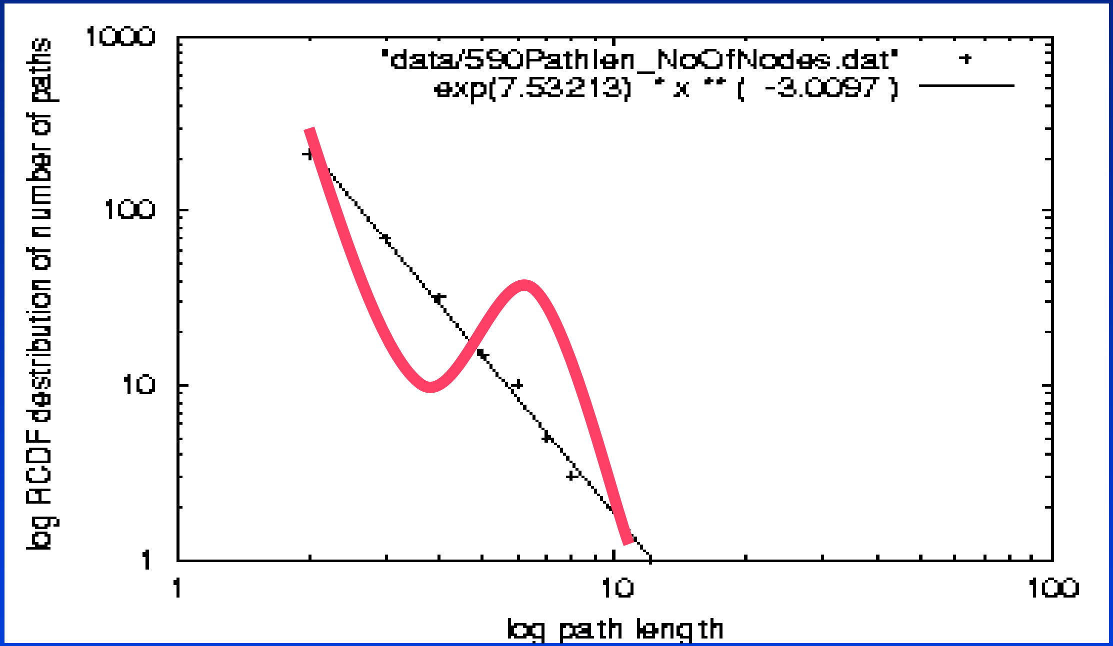

The Failure of the Donut Model

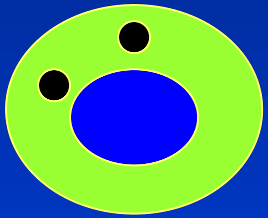

Path Length

- **All alternate paths go through the same direction**
- **No shortcuts or loop-arounds**

Roadmap

- Identify a Hierarchy
  - Defining the Importance of a node
- Present topological properties
- Present the jellyfish model
- Why is the jellyfish a good model?
- Conclusions

Defining a Hierarchy Recursively

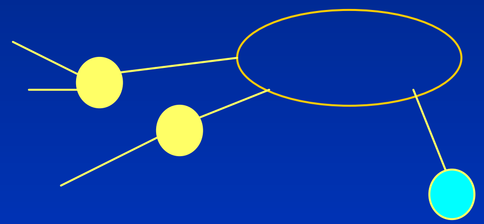

Define the core:

- Maximal clique of highest degree node

Define the Layers:

- All nodes adjacent to previous layer

Define the Shells:

- A layer without its one-degree nodes

The Internet Topology as a Jellyfish

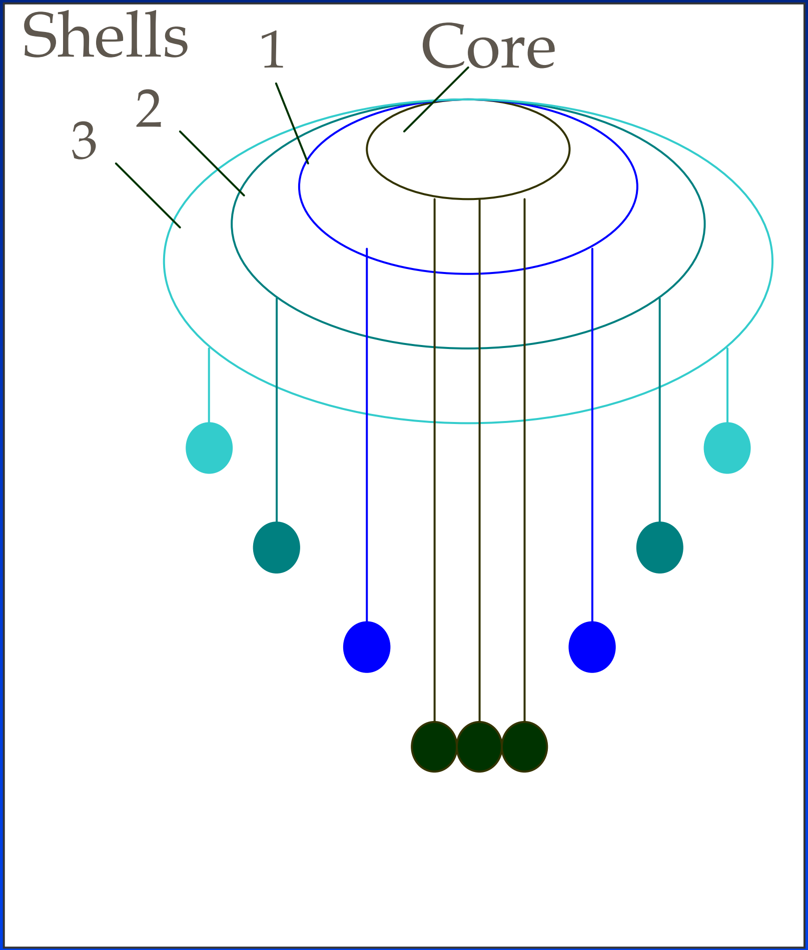

- Core: **High-degree clique**
- Shell: **adjacent nodes of previous shell, except 1-degree nodes**
- 1-degree nodes: **shown hanging**
- **The denser the 1-degree node population the longer the stem**

The Hierarchy: The Model Respects the Node Importance

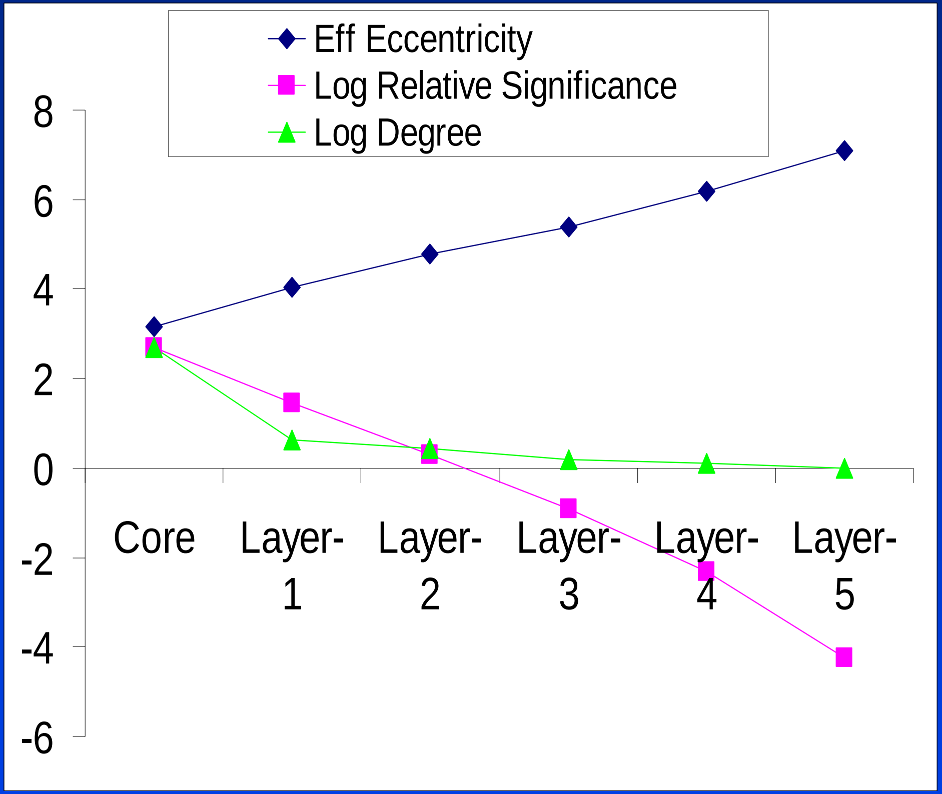

- The importance of nodes decreases as we move away from the core
- The effective eccentricity decreases by one in each layer

(see paper for details)

The Evolution of the Jellyfish

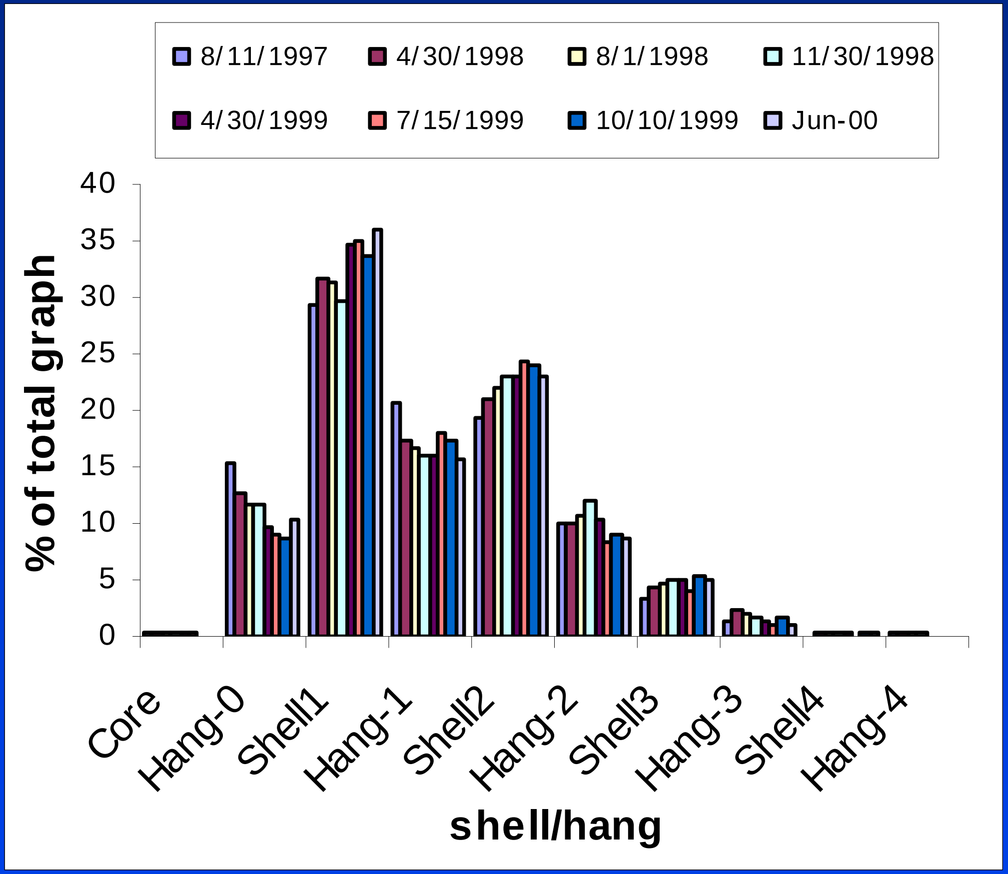

- The structure of the jellyfish has not changed much in 1997-2000
- Nodes become more connected:

  Small increase in shells and decrease in hanging nodes

The Diameter Remains Constant

Percentage of nodes reached

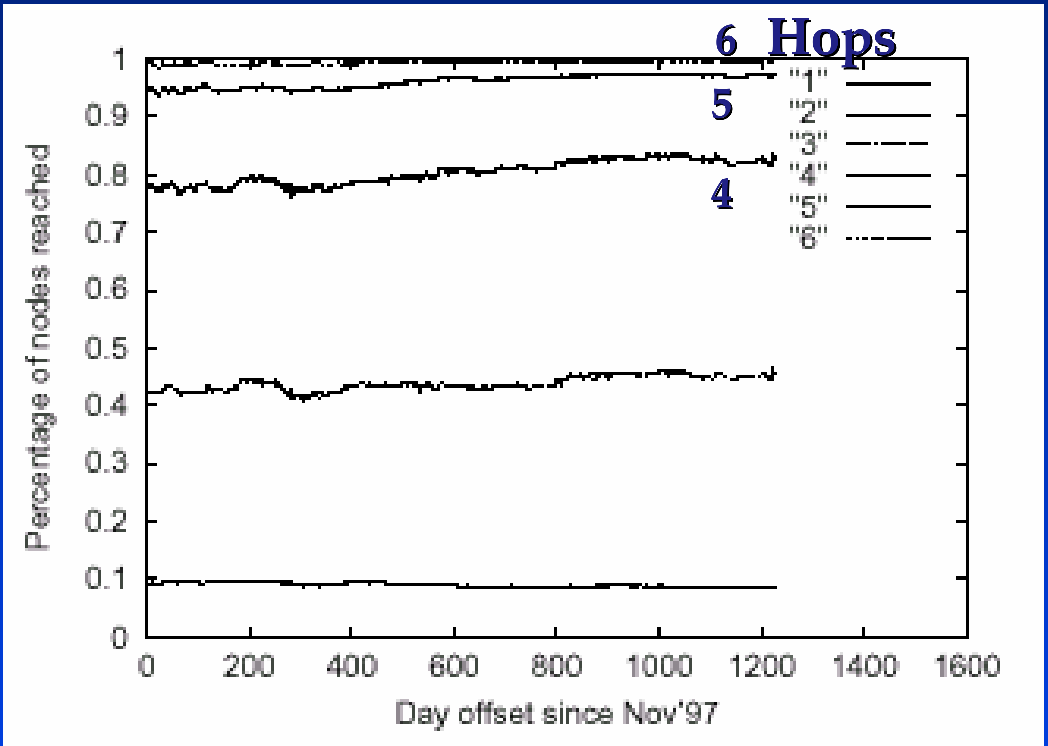

Time

- 6 hops reach approximately 98% of the network!
- The jellyfish diameter remains the same

Theory Supports the Jellyfish!

- **A surprising theoretical result [Reittu Norros 03]**
  - A network with powerlaw degree -> jellyfish
- **Assume degree powerlaw and random connections**
  - The network will have a clique of high degree nodes
  - The diameter of the network is O(log logN)!
- **In total aggrement with our observations**

Why Is The Jellyfish a Good Model?

It's cute, in addition…

The Jellyfish Captures Many Properties

- **The network is compact:**
  - 99% of pairs of nodes are within 6 hops
- **There exists a highly connected center**
  - Clique of high degree nodes
- **There exists a loose hierarchy:**
  - Nodes far from the center are less important
- **One-degree nodes are scattered everywhere**
- **The network has the tendency to be one large connected component**

And It Looks Like A Jellyfish…

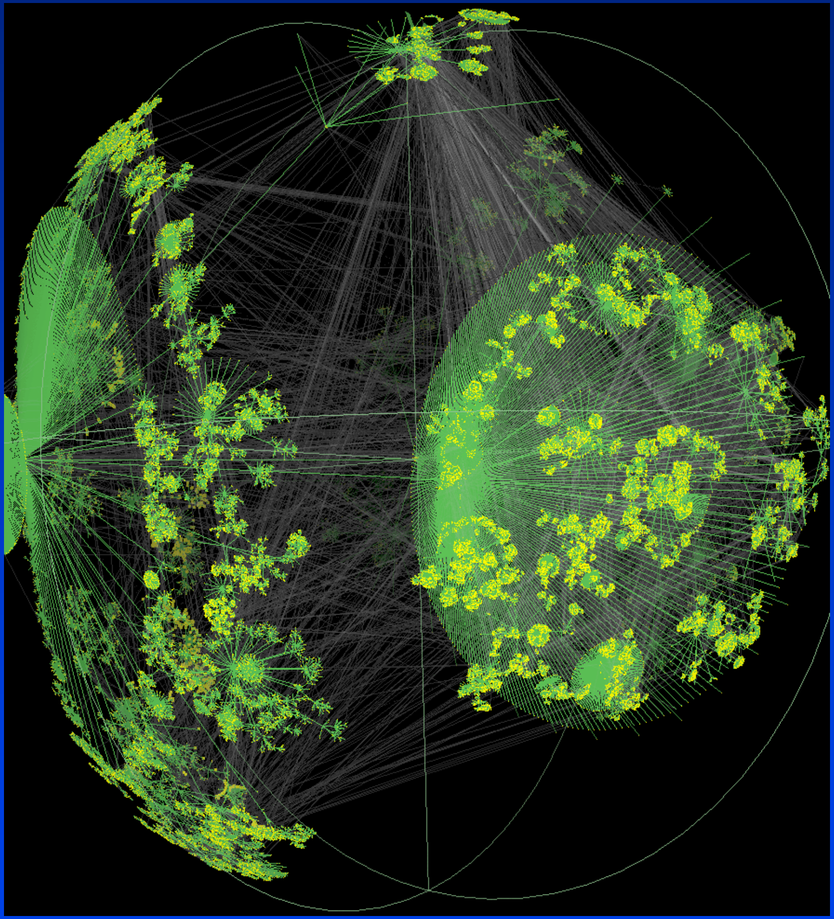

- Independent Observation
- Router Level Topology
- Produced by CAIDA

Conclusions

- **We model the Internet as a jellyfish**
- **The jellyfish represents graphically several topological properties**
  - Network is compact
  - We can identify a center
  - We can define a loose hierarchy
  - The network tends to be one connected component
- **Theoretical results support our observations**

http://www.cs.ucr.edu/~michalis/

My Other Research Interests

1. **Characterize and model network behavior:**
   - Poisson and Long Range Dependence [GI 02 - Globecom]
2. **Model and simulate the Internet topology**
   - Identify structure and hierarchy [GI 01] [COMNET*]
3. **Model and simulate BGP**
   - Large scale simulations (10,000 nodes) [GI 02]
4. **Wireless networks**
   - Improving TCP over ad hoc networks [Globecom 02]
5. **Multicast: supporting scalability and QoS (Cui Gerla)**
   - Efficient management through overlay trees [Globecom 02]+
6. **A novel network layer: DART (PeerNet)** [IPTPS 03]

Appendix:
Latest News on Power-laws

- **The Internet topology can be described by power-laws** **[Faloutsos x 3, SIGCOMM'99]**
- **The power-laws are here to stay**
  - Appear consistently over five years
  - Even with newer more complete data [Infocom'02]

Powerlaw: Degree Exponent D

RouteViews - NLANR Data

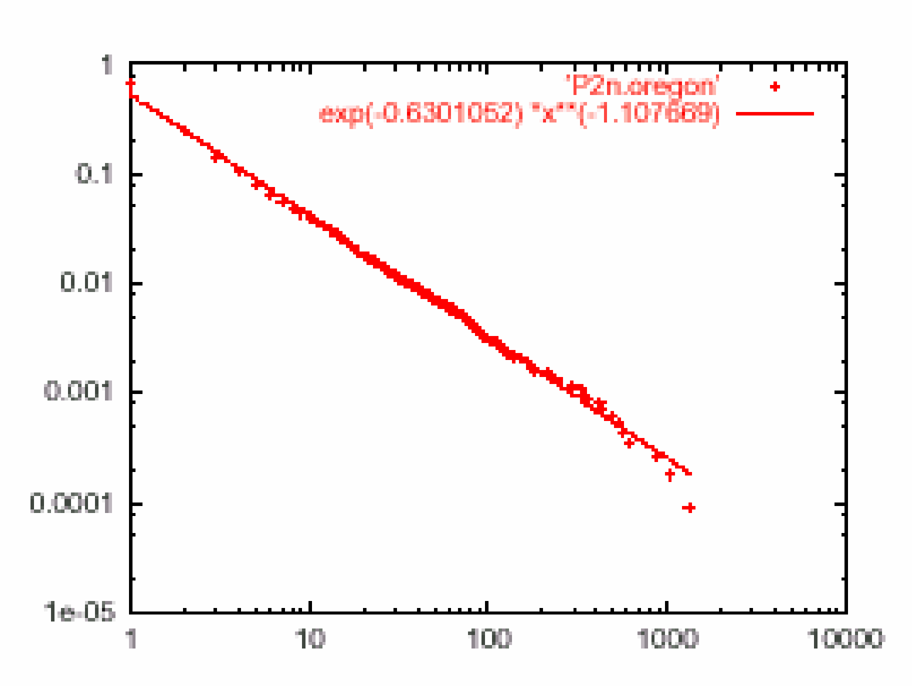

Newer More Complete AS graph

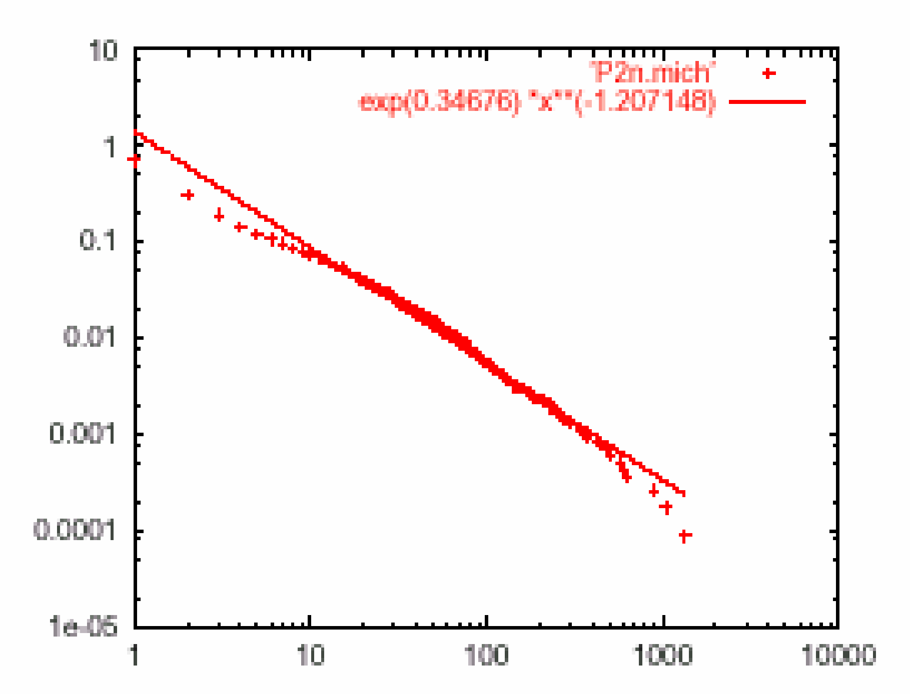

- **Degree distribution of nodes: CCDF**
- **It holds even for the more complete graph: 99%**

Thank you!

http://www.cs.ucr.edu/~michalis/

Eigenvalues of the Topology

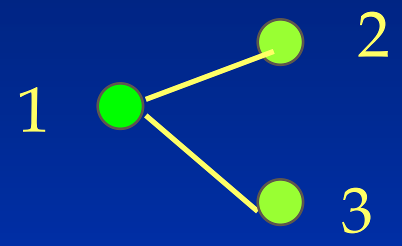

$$A = \begin{vmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{vmatrix}$$

- **Let *A* be the adjacency matrix of graph**
- **The eigenvalue λ is real number s.t.:**
  - $A\,\underline{v} = \lambda\,\underline{v}$, where $\underline{v}$ some vector
- **Eigenvalues are strongly related to topological properties**
- **More details in Part B**

Power-law: Eigen Exponent *E*

Eigenvalue

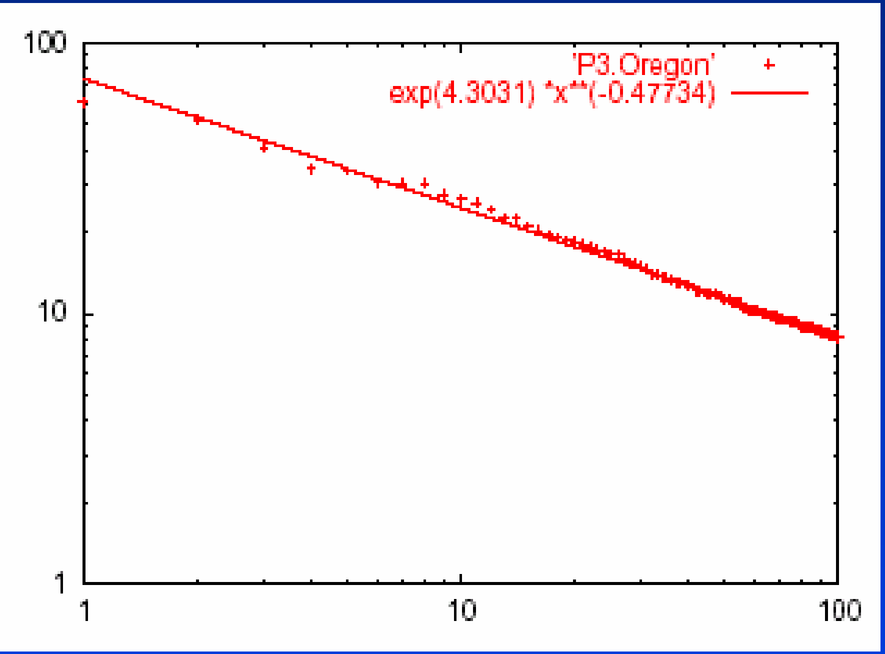

Rank of decreasing eigenvalue

Exponent = slope

E = -0.48

May 2001

- **Find the eigenvalues of the adjacency matrix**
- **Eigenvalues in decreasing order (first 100)**

Surprising Result!

- **Exponent E is half of exponent D**
- **Theorem: Given a graph with relatively large degrees d_i then with high probability:**
  - Eigenvalue $\lambda_i = \sqrt{d_i}$, where i rank of decreasing order
- **Thus, if we compare the slope of the plot the eigenvalues and the degrees:**
  - $\log \lambda_i = 0.5 \log d_i$

    [Fabrikant, Koutsoupias, Papadimiitriou in STOC'01]

    [Mihail Papadimitriou Random 02]

Time Evolution of The Topology

- Powerlaws are here to stay
- Degree distribution slope is invariant
- Network becomes denser
- The rich get richer phenomenon

The Number of ASes in Time

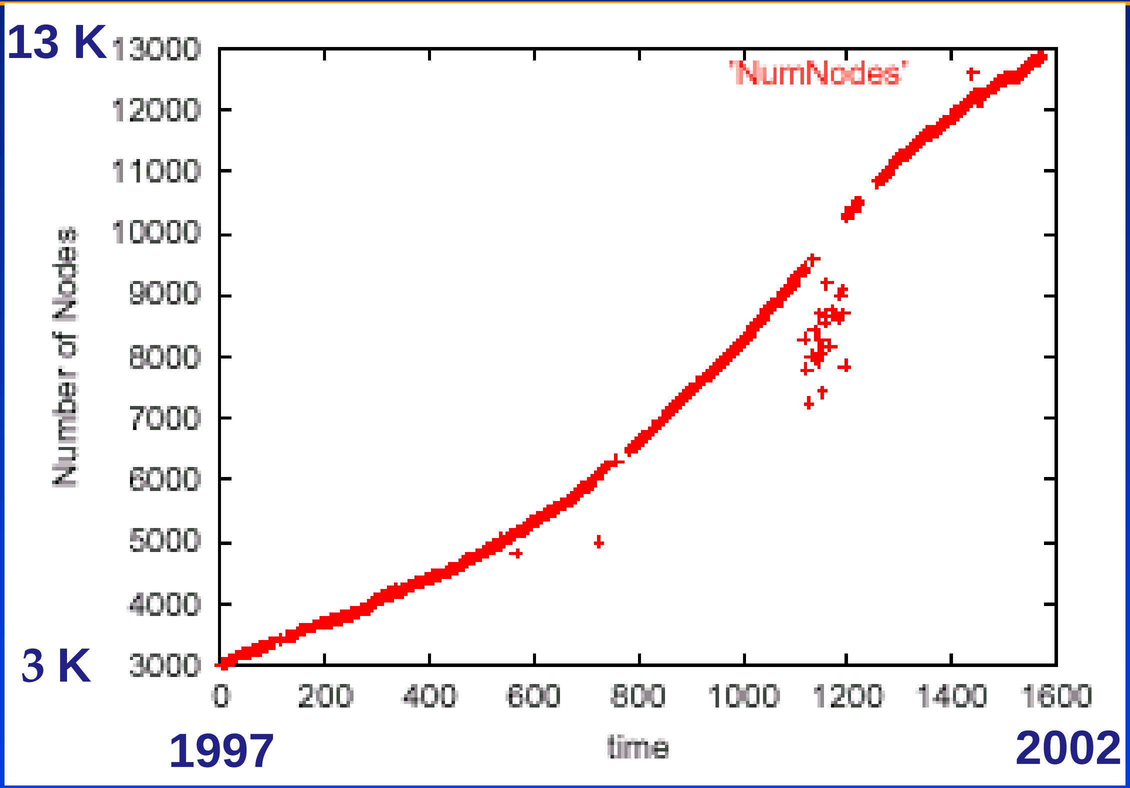

- The number of AS doubled in two years
- Growth seems to slow down!

Degree Distribution Did Not Change!

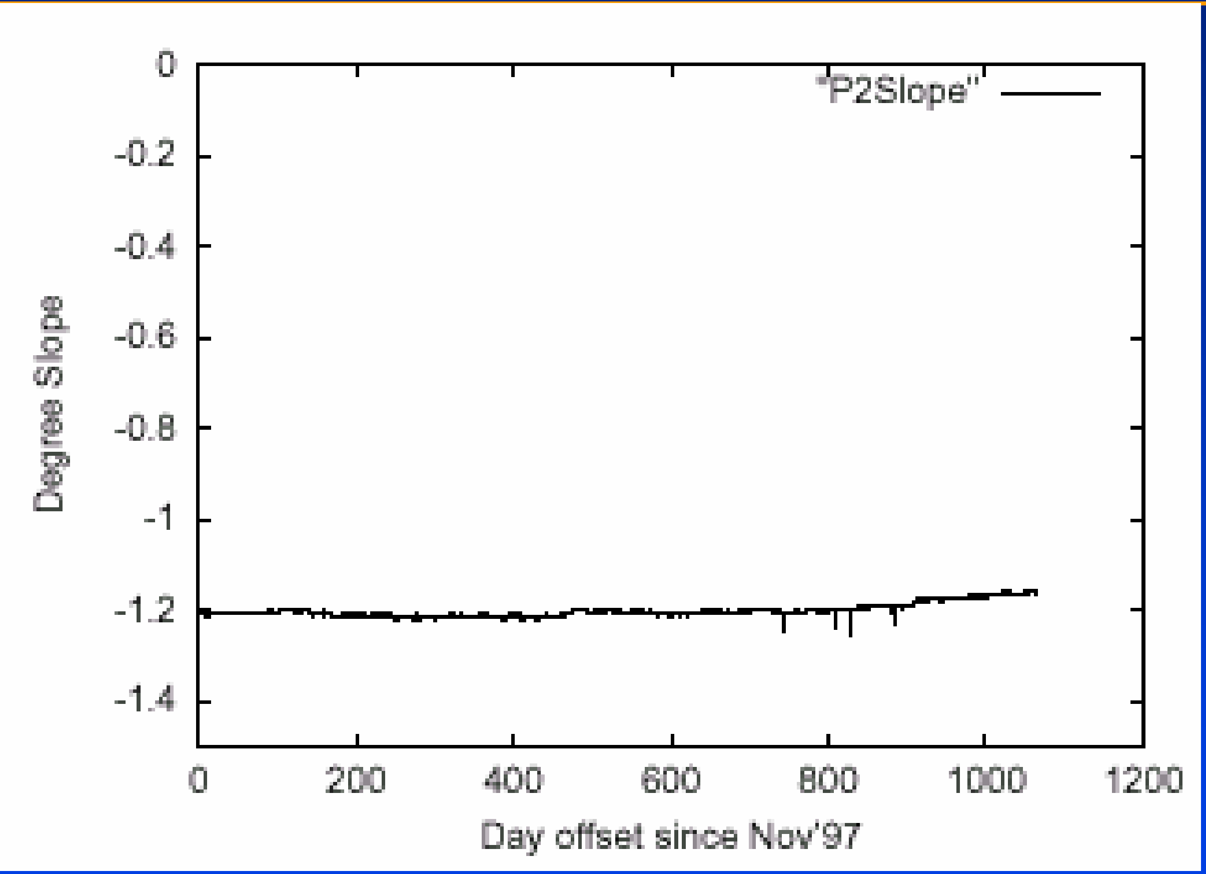

- **Slope is practically constant for over 3 years**

The Topology Becomes Denser!

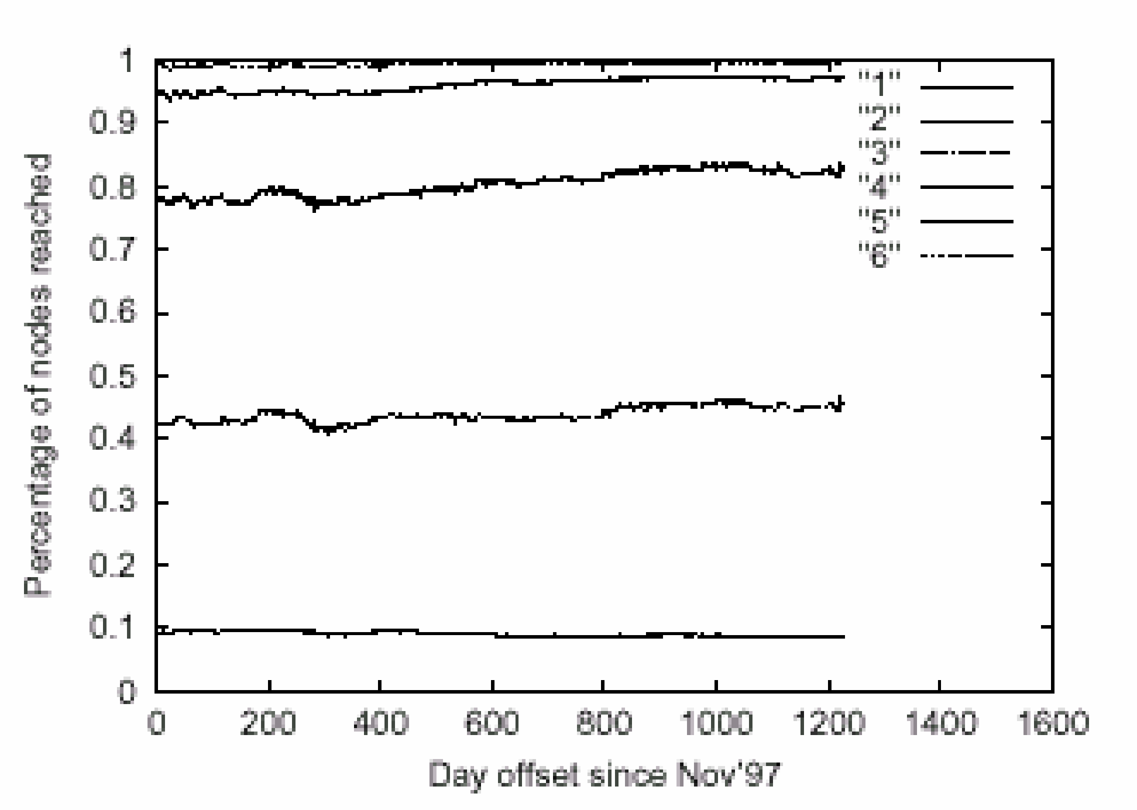

Recall six degrees of separation

- 6 hops reach approximately 98% of the network!
- Denser: 6 hops reach more nodes

The Rich Get Richer

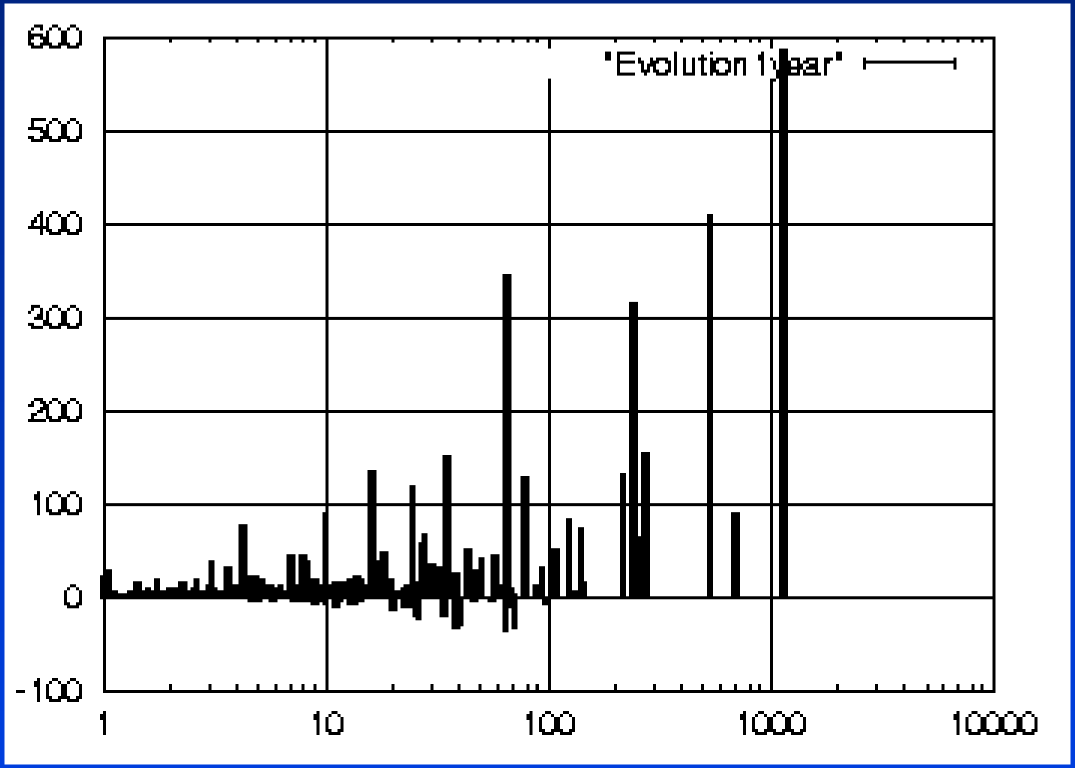

- The increase of the degree versus the initial degree
- New connections prefer “highly connected nodes”

That's it!

Thank you!

http://www.cs.ucr.edu/~michalis/

I. Power-law: rank exponent *R*

degree

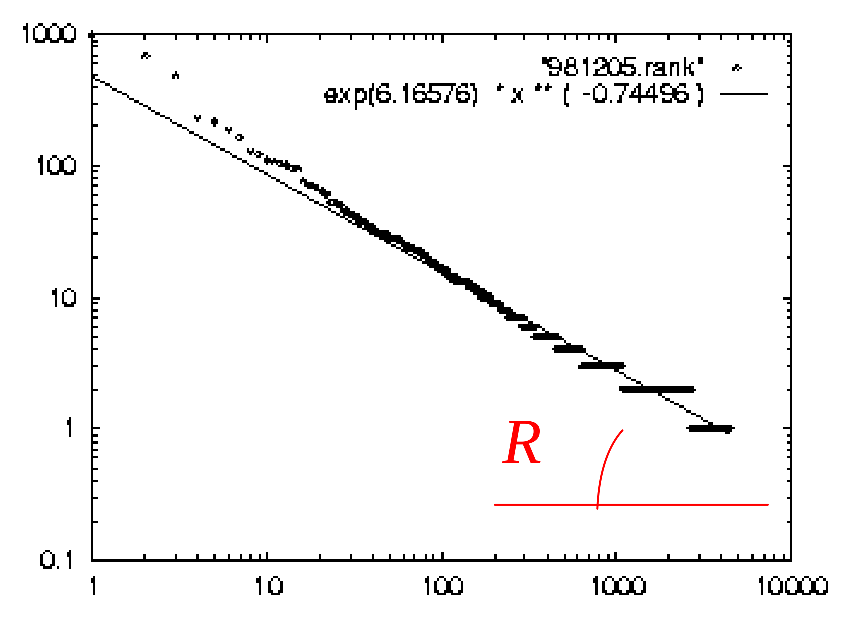

Exponent = slope

R = -0.74

Dec'98

Rank: nodes in decreasing degree order

- **The plot is a line in log-log scale**

[Faloutsos, Faloutsos and Faloutsos SIGCOMM'99]

I. Estimations Using With Rank Exponent *R*

Lemma:

Given the nodes *N*, and an estimate for the rank exponent *R*, we predict the edges E:

$$E = \frac{1}{2(R+1)} \cdot \left(1 - \frac{1}{N^{R+1}}\right) \cdot N$$

University of California Riverside

Some Current Results

- **Measuring the performance of real-time applications**
  - E2e performance is asymmetric (by 10)
- **Estimating Long Range Dependence**
  - No definitive estimating method exists
  - SELFIS software tool for performance analysis
- **A study of BGP routing robustness**
  - Persistence and prevalence of paths
  - Paths are fairly robust, but there is a lot of “noise” too
  - A data repository: 107Gb, 1 billion BGP paths

Measuring Real-time Performance

- "Can the Internet support VoIP <u>now</u>?"
- We conduct globe-wide experiments
  - UCR, CMU, Japan, Australia, Greece
- Experimental set-up
  - Approx. 6 4Kbit/sec sending rate
  - Small packet sizes every 20, 30, 40, 50 msec, 1 sec

How Many Distinct Paths Does an IP Prefix Use?

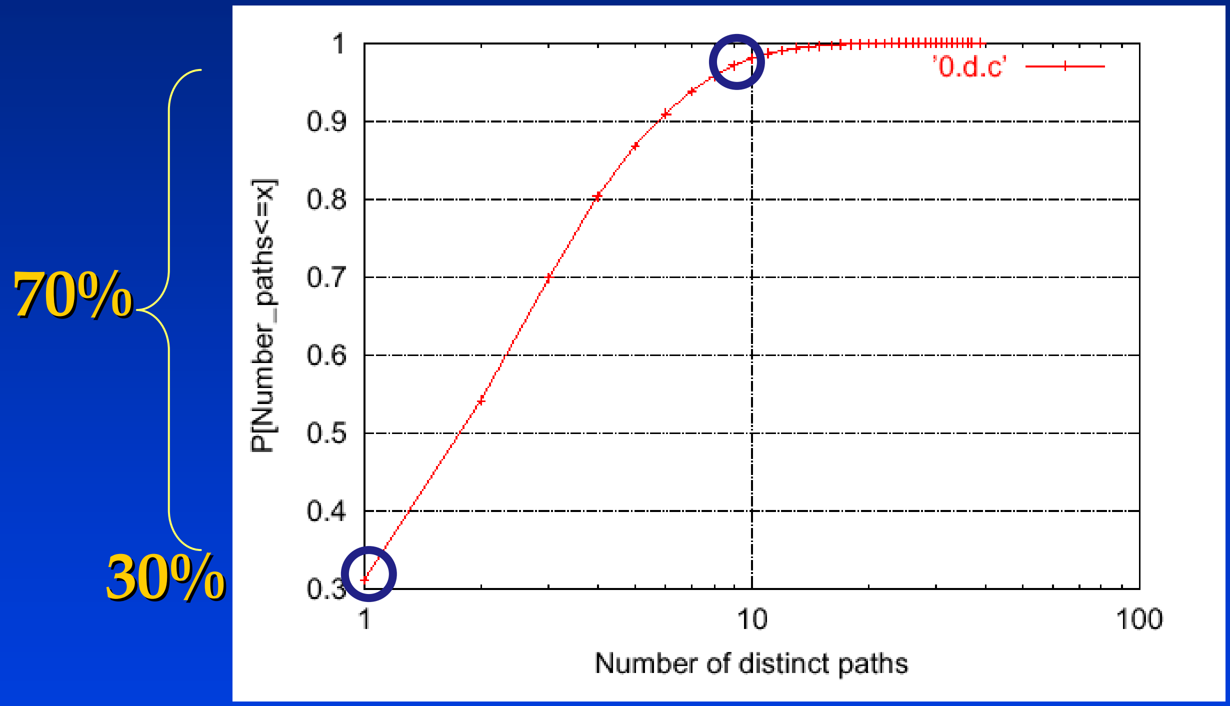

- Almost 70% of the IP prefixes have 2-10 distinct paths
- 30% of IP prefixes have only one path

For More Information

www.cs.ucr.edu/~michalis/

Routing Is Persistent

70%

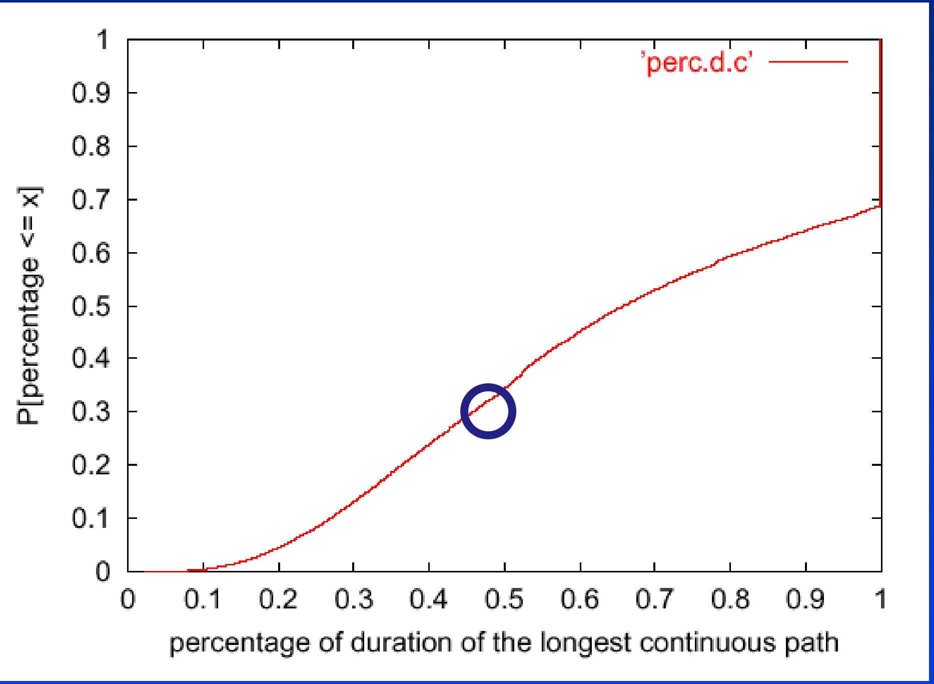

70% of prefixes use one path continuously for 50% of their time!

CDF of the relative duration of the most persistent path

Measurements: The Death of the Symmetry Assumption

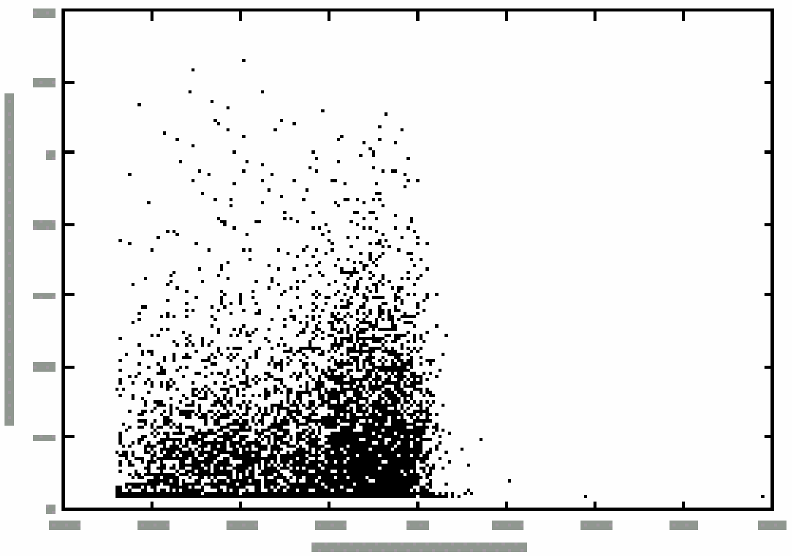

Fig. 1. CMU dataset with 40-msec sending rate and 320-byte packet size.

- **One-way delay:**
  Forward can be 10 times higher than backward delay

Characterizing Network Behavior with Long Range Dependence

- **LRD captures the "memory" of the behavior**
- **It is quantified by a single scalar number**
- **LRD appears in many aspects of networks**
  - Traffic load, arrival times, delays, packet loss
- **Open Question: what does it really tell us?**

<u>PROBLEM:</u> We do not know how to calculate LRD!

- Many estimators with conflicting estimates
- No systematic approach

The Intuition Behind LRD

- Capturing the "dependency" of the signal to its previous values

- **White Noise**

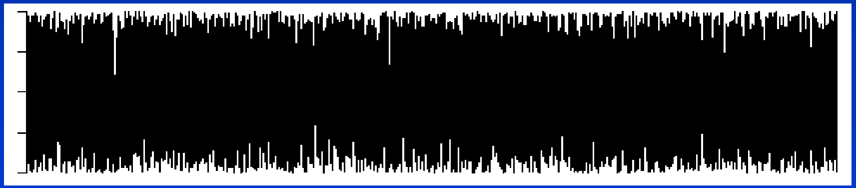

- **Pink Noise**

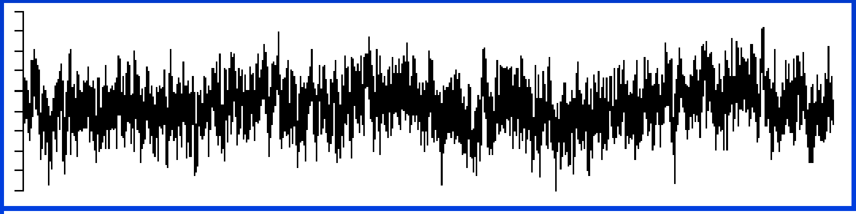

- **Brownian Noise**

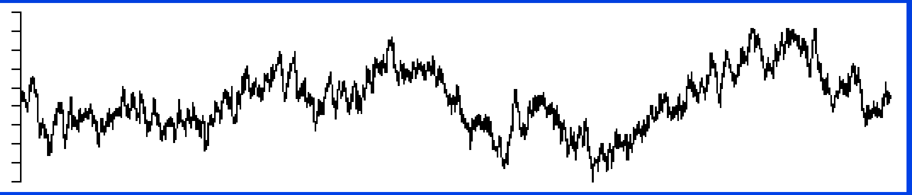

Idea: Reverse Engineering LRD

- **Develop a library of behaviors to know data**
- **Three Series of Tests for the Estimators**

1. **Evaluating the accuracy of the estimators**
   - Synthetic Fractional Gaussian Noise (FGN)
2. **Deceiving the estimators with non-LRD data**
   - Periodicity, Noise, Trend
3. **Applying the estimators in real data**
   - Characterizing delay and packet loss

BGP Routing Analysis

- **Overarching Goal:**
  - Develop a realistic detailed model for large scale realistic simulations
- **Now: A study of BGP routing robustness**
  - Persistence and prevalence of paths
  - Stability of advertisements
- **Next step:**
  - Study the customer-provider relationships

Using Massive BGP Routing Data

- **We use data from NLANR for almost 3 years**
  - Late 1997 to early 2001
- **Daily snapshots of BGP routing tables**
- **Created a database to facilitate path queries**
  - 107Gb of data, 1 billion BGP paths

Overview of Results for BGP Routing

Stable and persistent routing with some "noise"

- 44% prefixes are advertised for < 30 days
- 50% prefixes have a dominant path 84% of time
- 35% of prefixes use one path continuously for 90% of their time!
- Significant path multiplicity due to traffic engin.

Graph Reduction Tools

Large

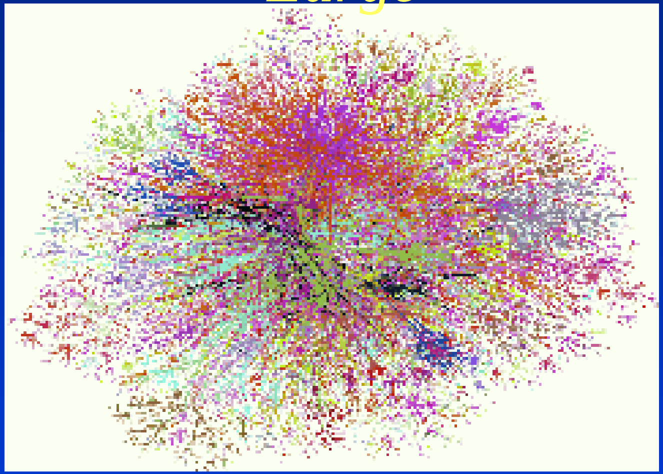

Reduce

Small

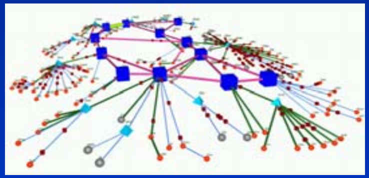

- **Reduce: large real graph to small realistic graph**
  - Achieve 70% reduction
- **Satisfy degree distribution, but increases diameter**

The Jellyfish Captures "Direction" of Connectivity

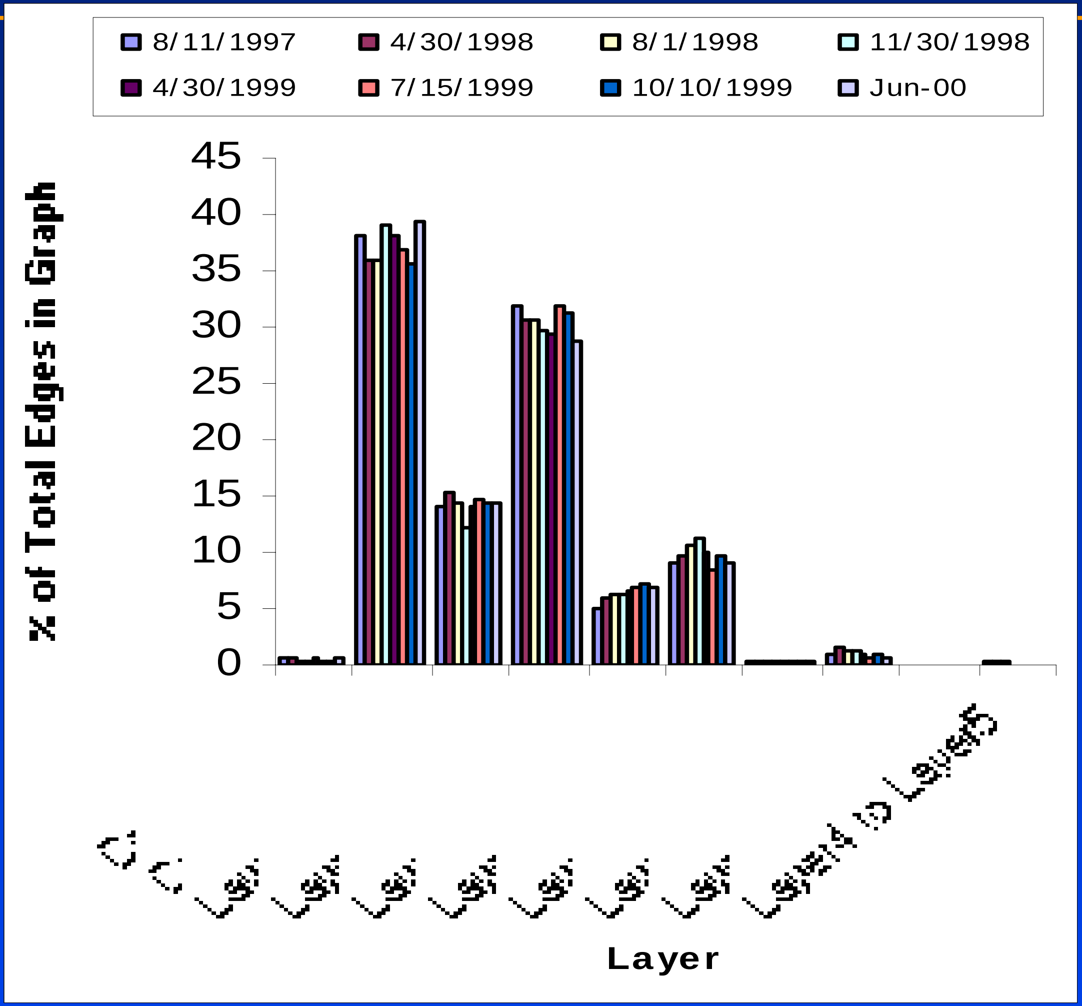

- Most edges are between layers 80%
- Less edges are within a layer 20%

The Model Respects the Node Importance

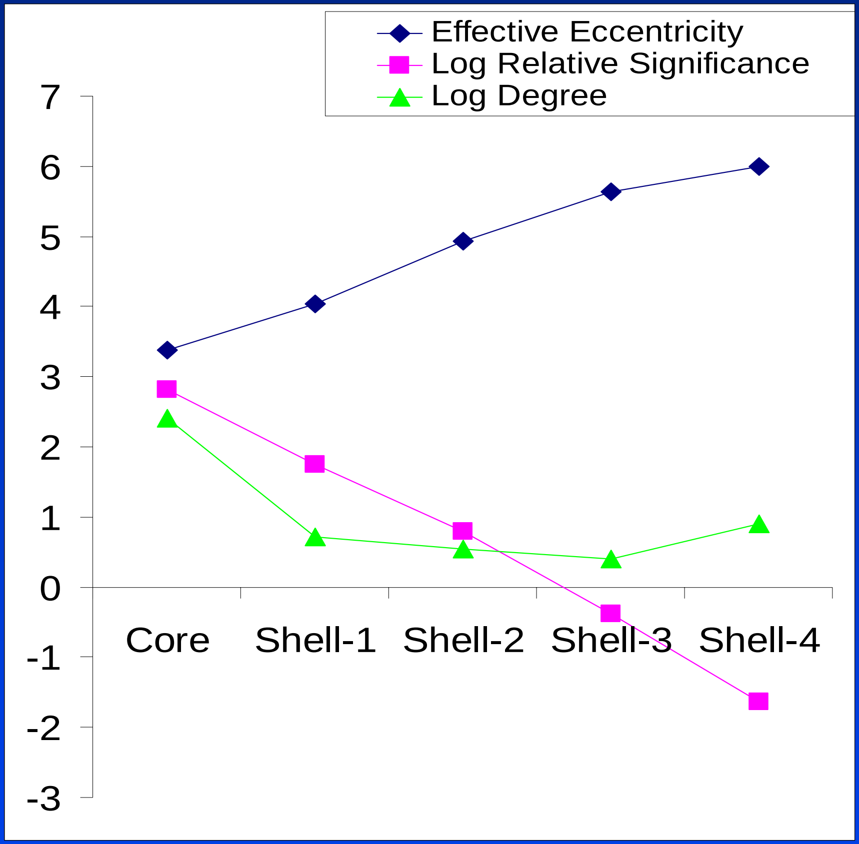

- The importance of each layer decreases as we move away from the core

Intuitive Models Are Useful

- **Cons:**
  - Danger of oversimplification
- **Pros:**
  - Memorable
  - Visualizable
  - Maximizing information/effort ratio
- **They can be very useful when exploring unknown territory**
  - Even disproving a wrong model is progress!