“(White) Lies, Damned Lies, and Statistics”
or
On measuring and modeling engineered networks: The Internet as a case study

Walter Willinger
AT&T Labs-Research
walter@research.att.com

Acknowledgments

Main Collaborators

- John Doyle (Caltech)
- David Alderson (Caltech)
- Lun Li (Caltech)

Contributions

- Reiko Tanaka (RIKEN, Japan)
- Matt Roughan (U. Adelaide, Australia)
- Steven Low (Caltech)
- Ramesh Govindan (USC)
- Stanislav Shalunov (Abilene)
- Heather Sherman (CENIC)

Basic Question

Do the available
Internet-related connectivity measurements
support the sort of claims that can be found in the
existing complex networks literature?

Key Issues

- Insist on a certain degree of data hygiene
- Insist on a certain level of statistical rigor
- Insist on taking model validation more serious

Internet connectivity structures are different at each layer

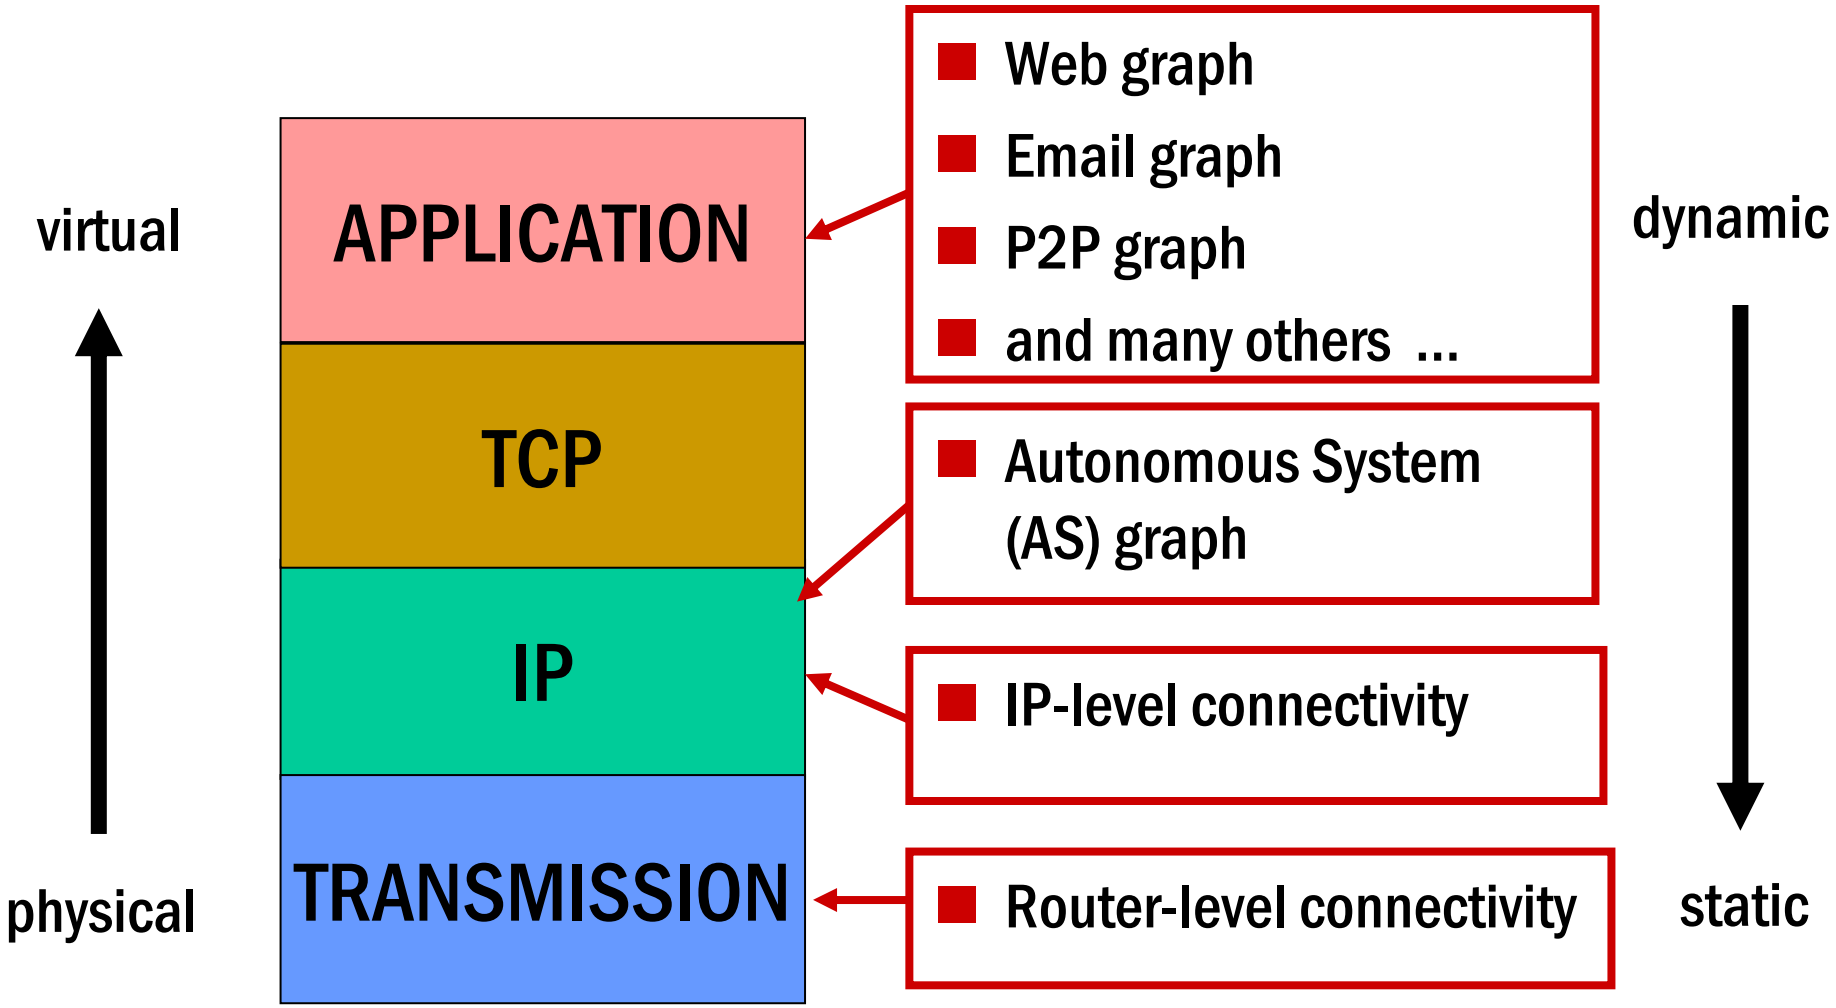

On Measuring Internet Connectivity

- No central agency/repository
- Economic incentive for ISPs to obscure network structure
- Direct inspection is typically not possible
- Based on measurement experiments, hacks
- Mismatch between what we want to measure and can measure

Specific Examples covered in this talk

- Physical infrastructure of an ISP
  - router-level topology
- Logical connectivity structures
  - AS-level topology
  - Web graph
  - P2P networks

On Measuring the Internet's Router-level Topology

- **`traceroute`** tool
  - Discovers compliant (i.e., IP) routers along path between selected network host computers

- Large-scale traceroute experiments
  - Pansiot and Grad (router-level map from around 1995)
  - Cheswick and Burch (mapping project 1997--)
  - Mercator (router-level maps from around 1999 by R. Govindan and H. Tangmunarunkit)
  - Skitter (ongoing mapping project by CAIDA folks)
  - Rocketfuel (state-of-the-art router-level maps of individual ISPs by UW folks)
  - Dimes (EU project)

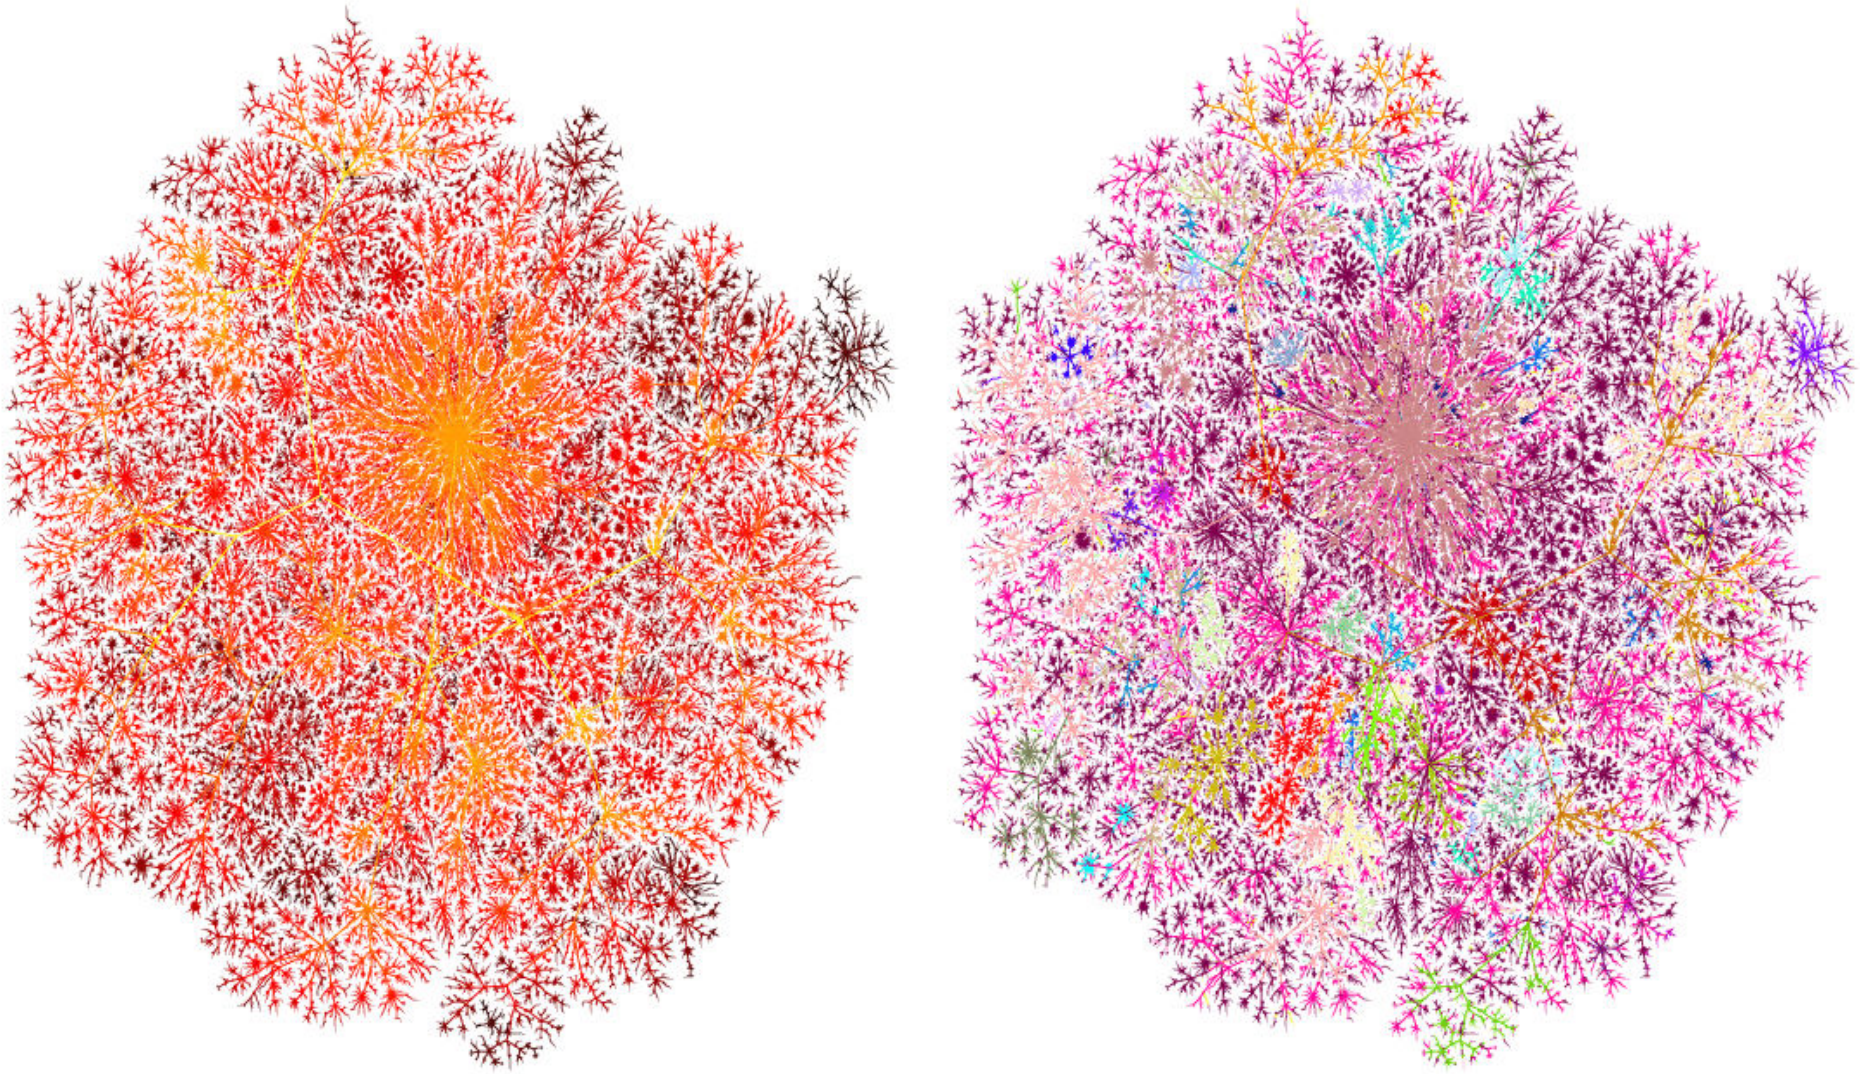

http://research.lumeta.com/ches/map/

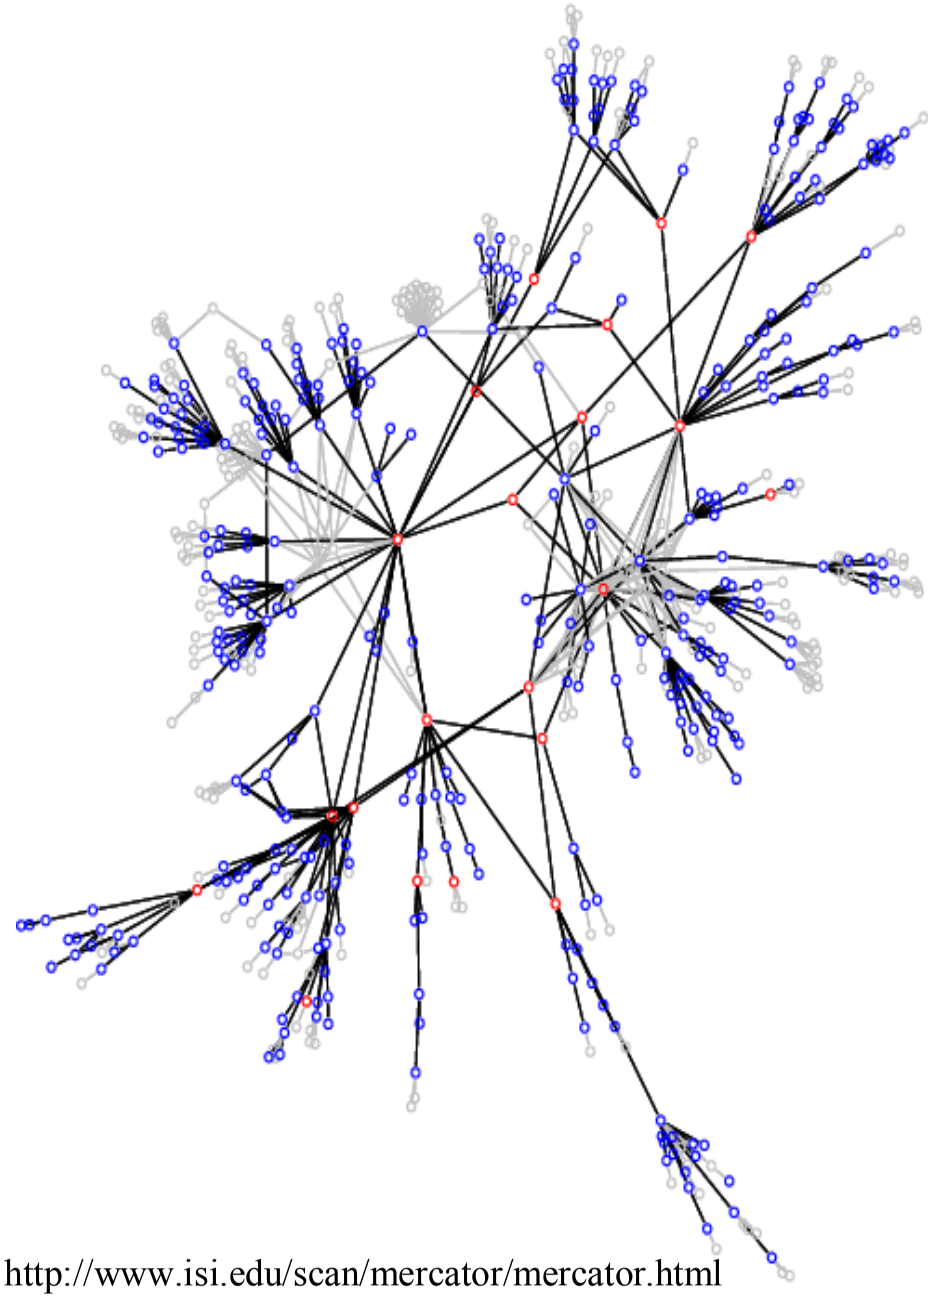

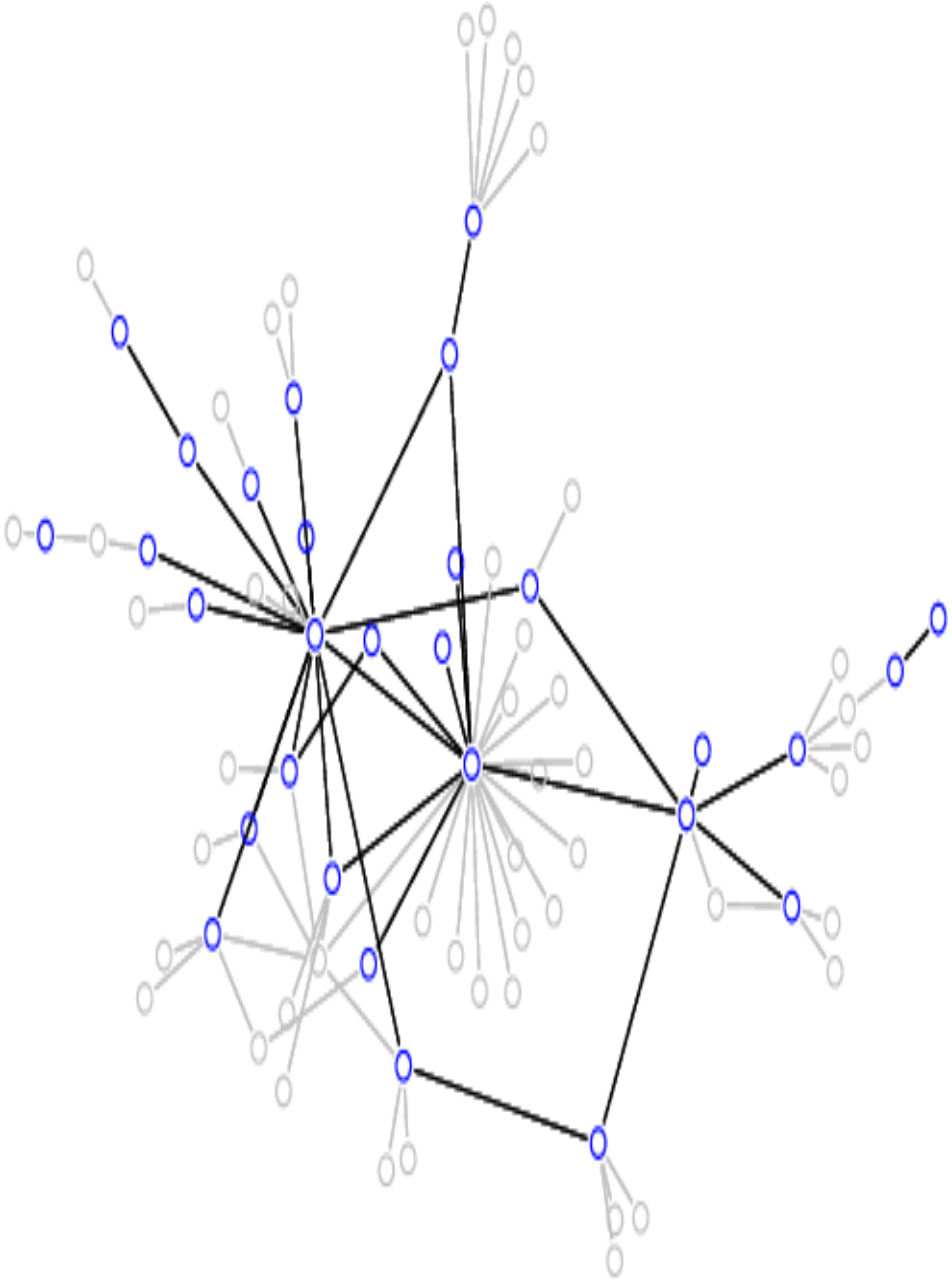

http://www.isi.edu/scan/mercator/mercator.html

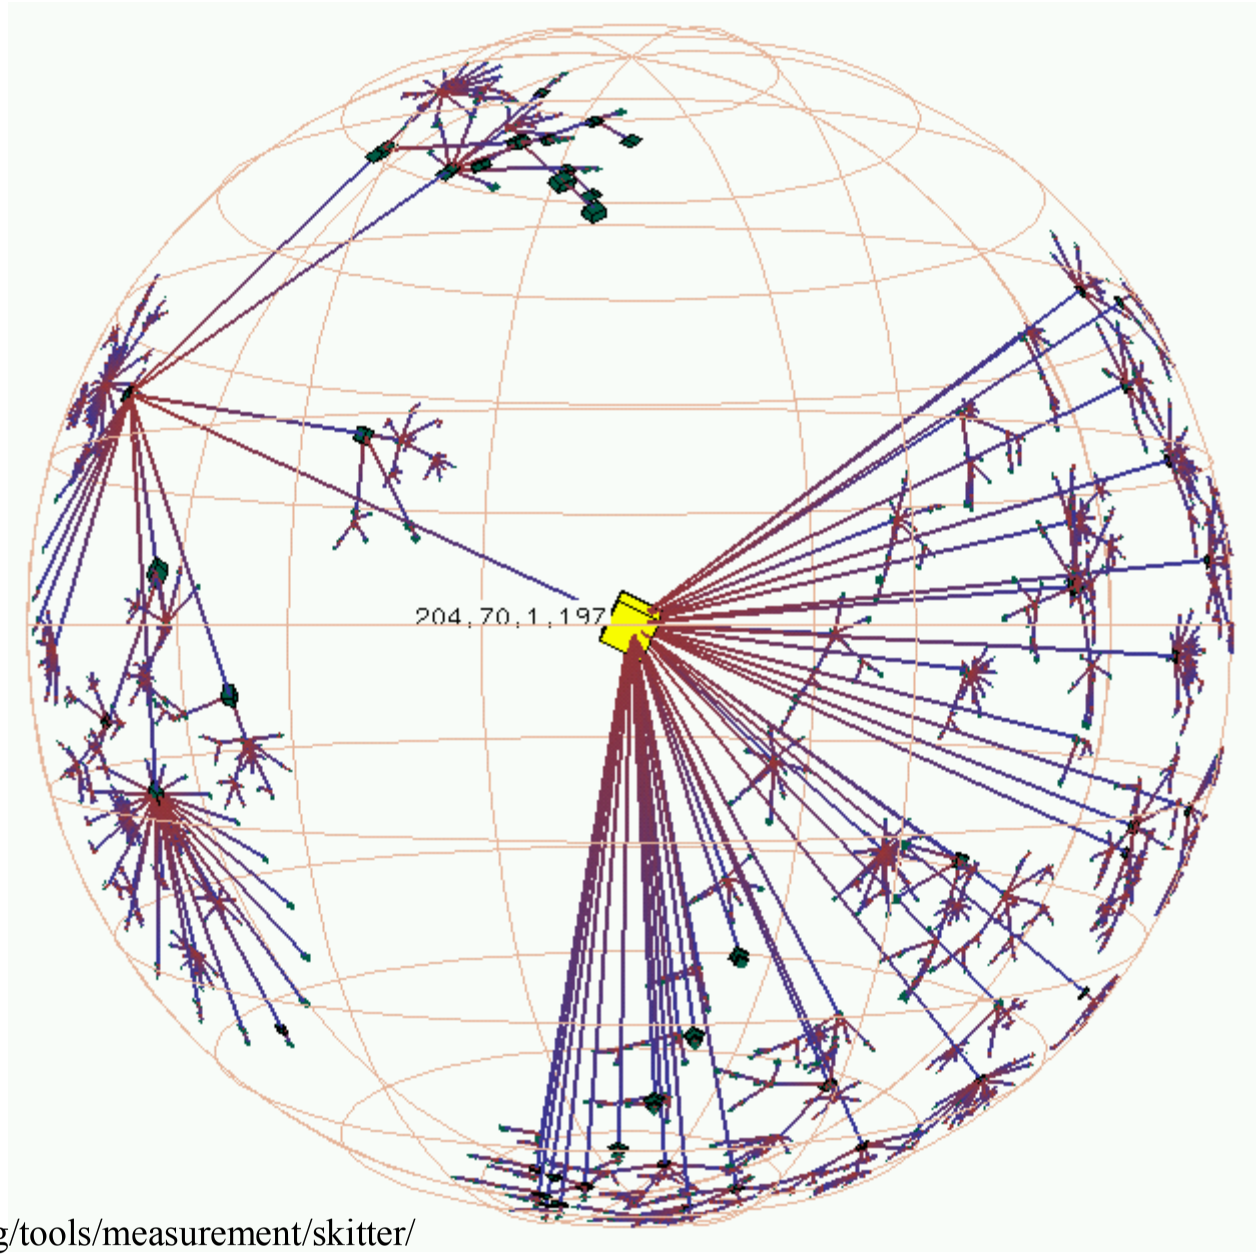

http://www.caida.org/tools/measurement/skitter/

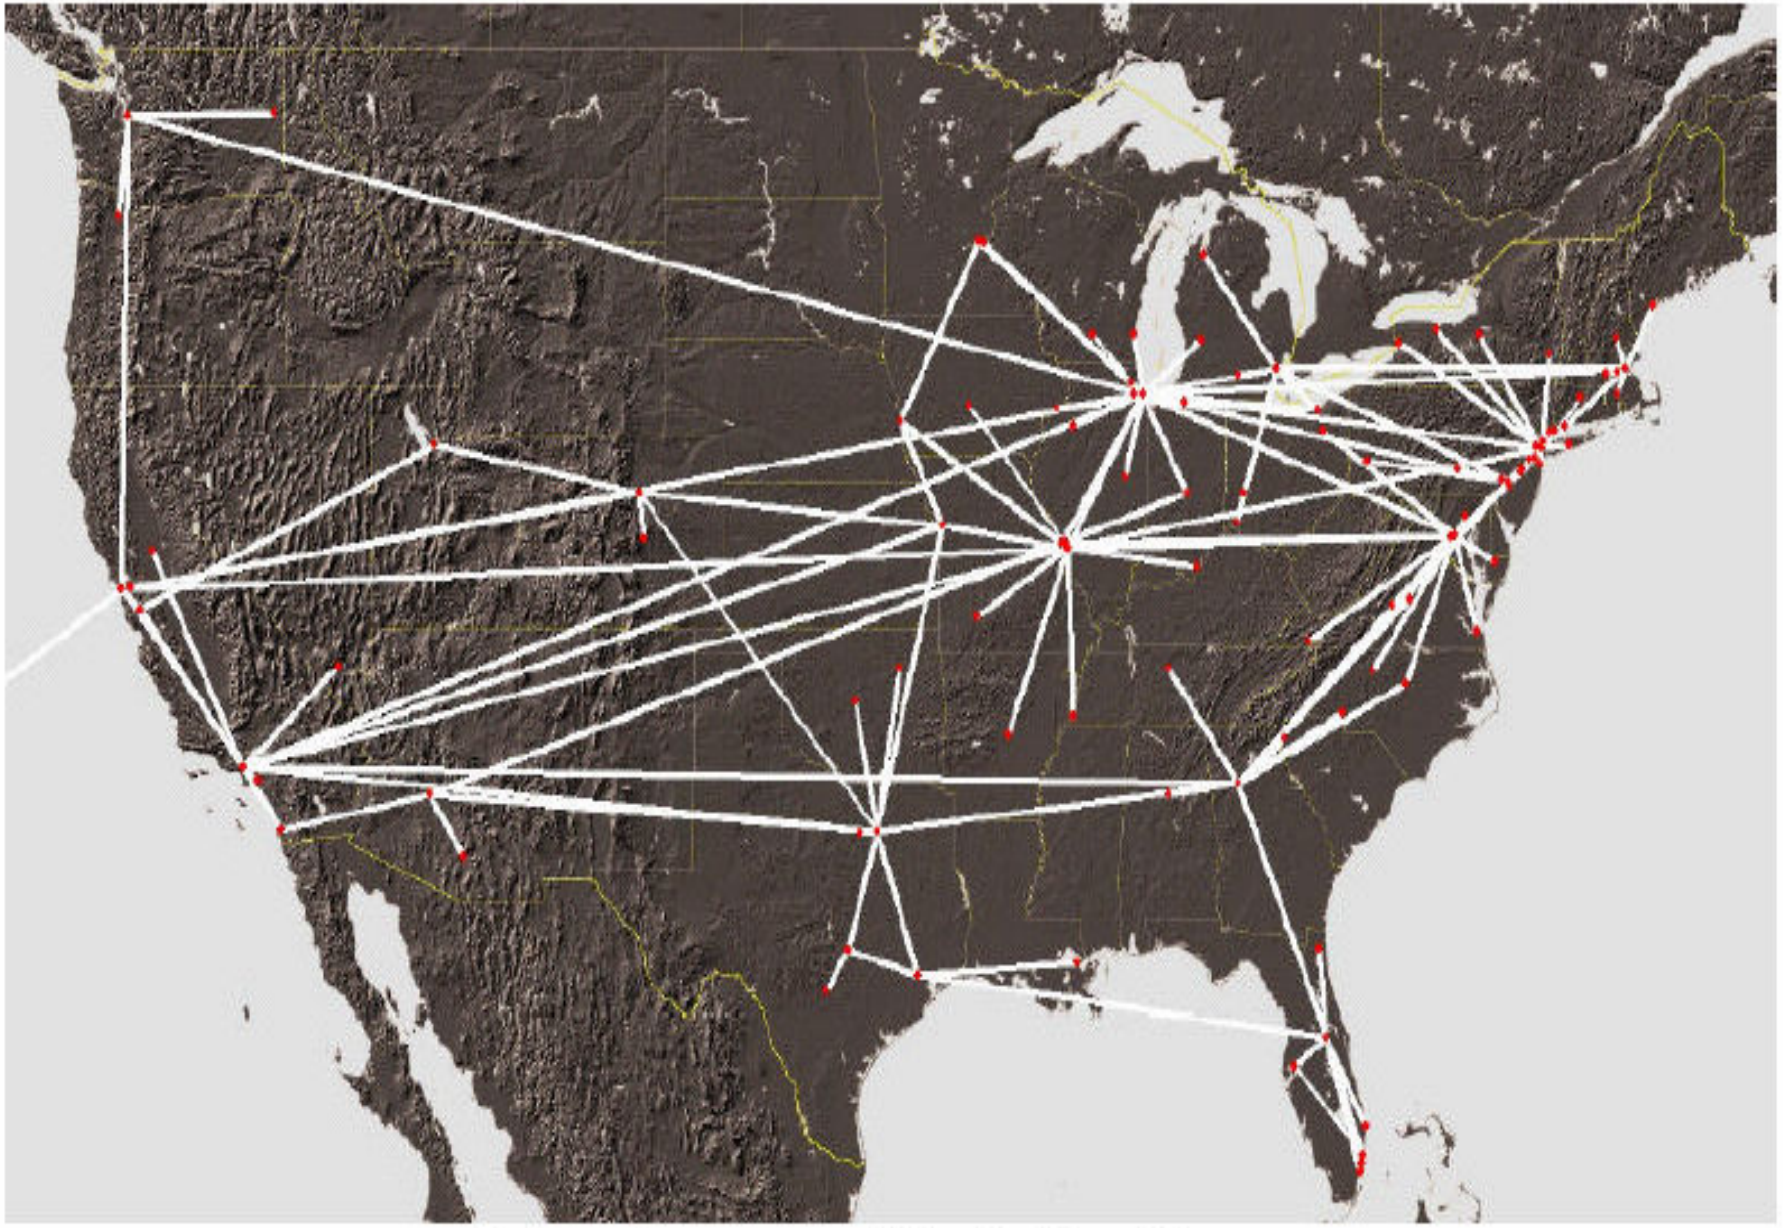

Background image courtesy JHU, applied physics labs

http://www.cs.washington.edu/research/networking/rocketfuel/bb

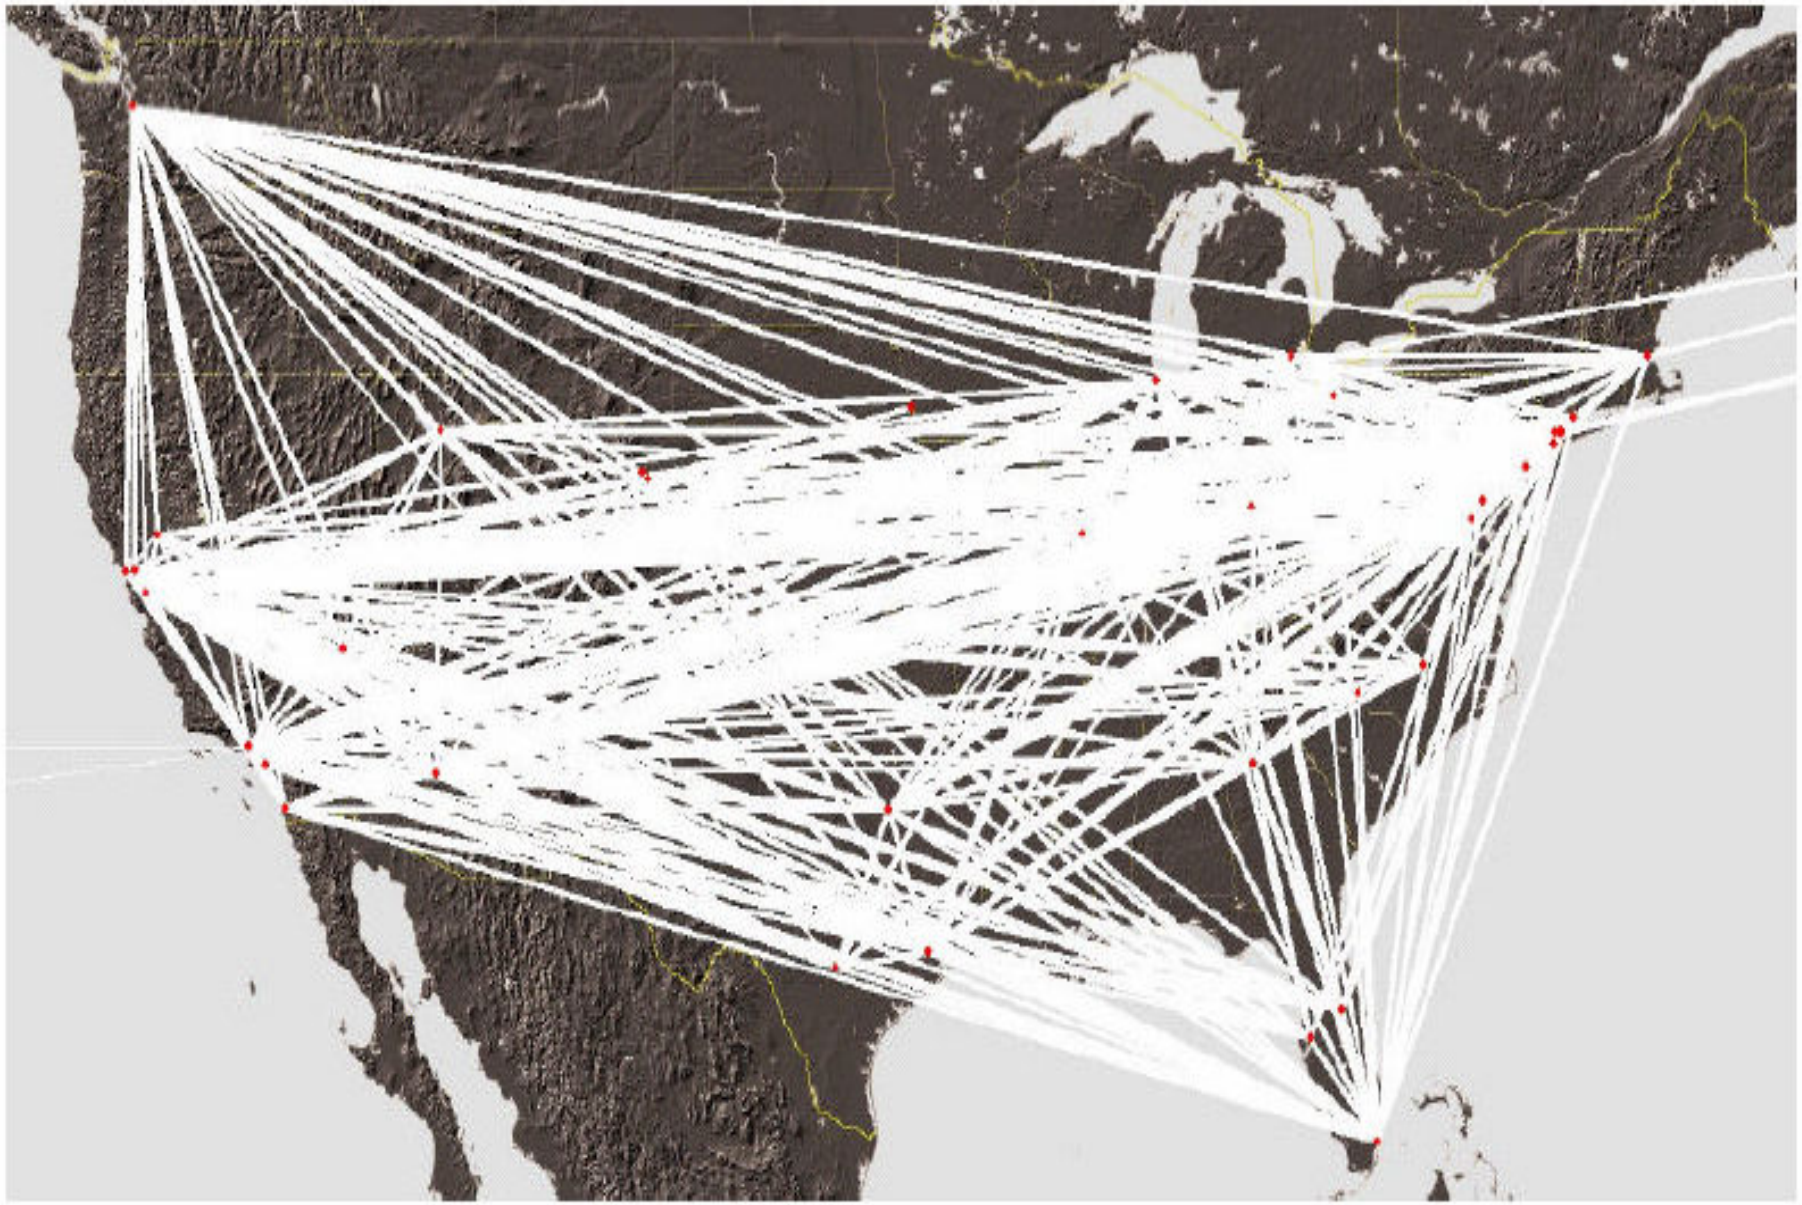

Background image courtesy JHU, applied physics labs

http://www.cs.washington.edu/research/networking/rocketfuel/

HOWEVER: Problems with existing measurements

- traceroute-based measurements are ambiguous
  - traceroute is strictly about IP-level connectivity
  - traceroute cannot distinguish between high connectivity nodes that are for real and that are fake and due to underlying Layer 2 (e.g., Ethernet, ATM) or Layer 2.5 technologies (e.g., MPLS)

The Internet: The Engineering Perspective

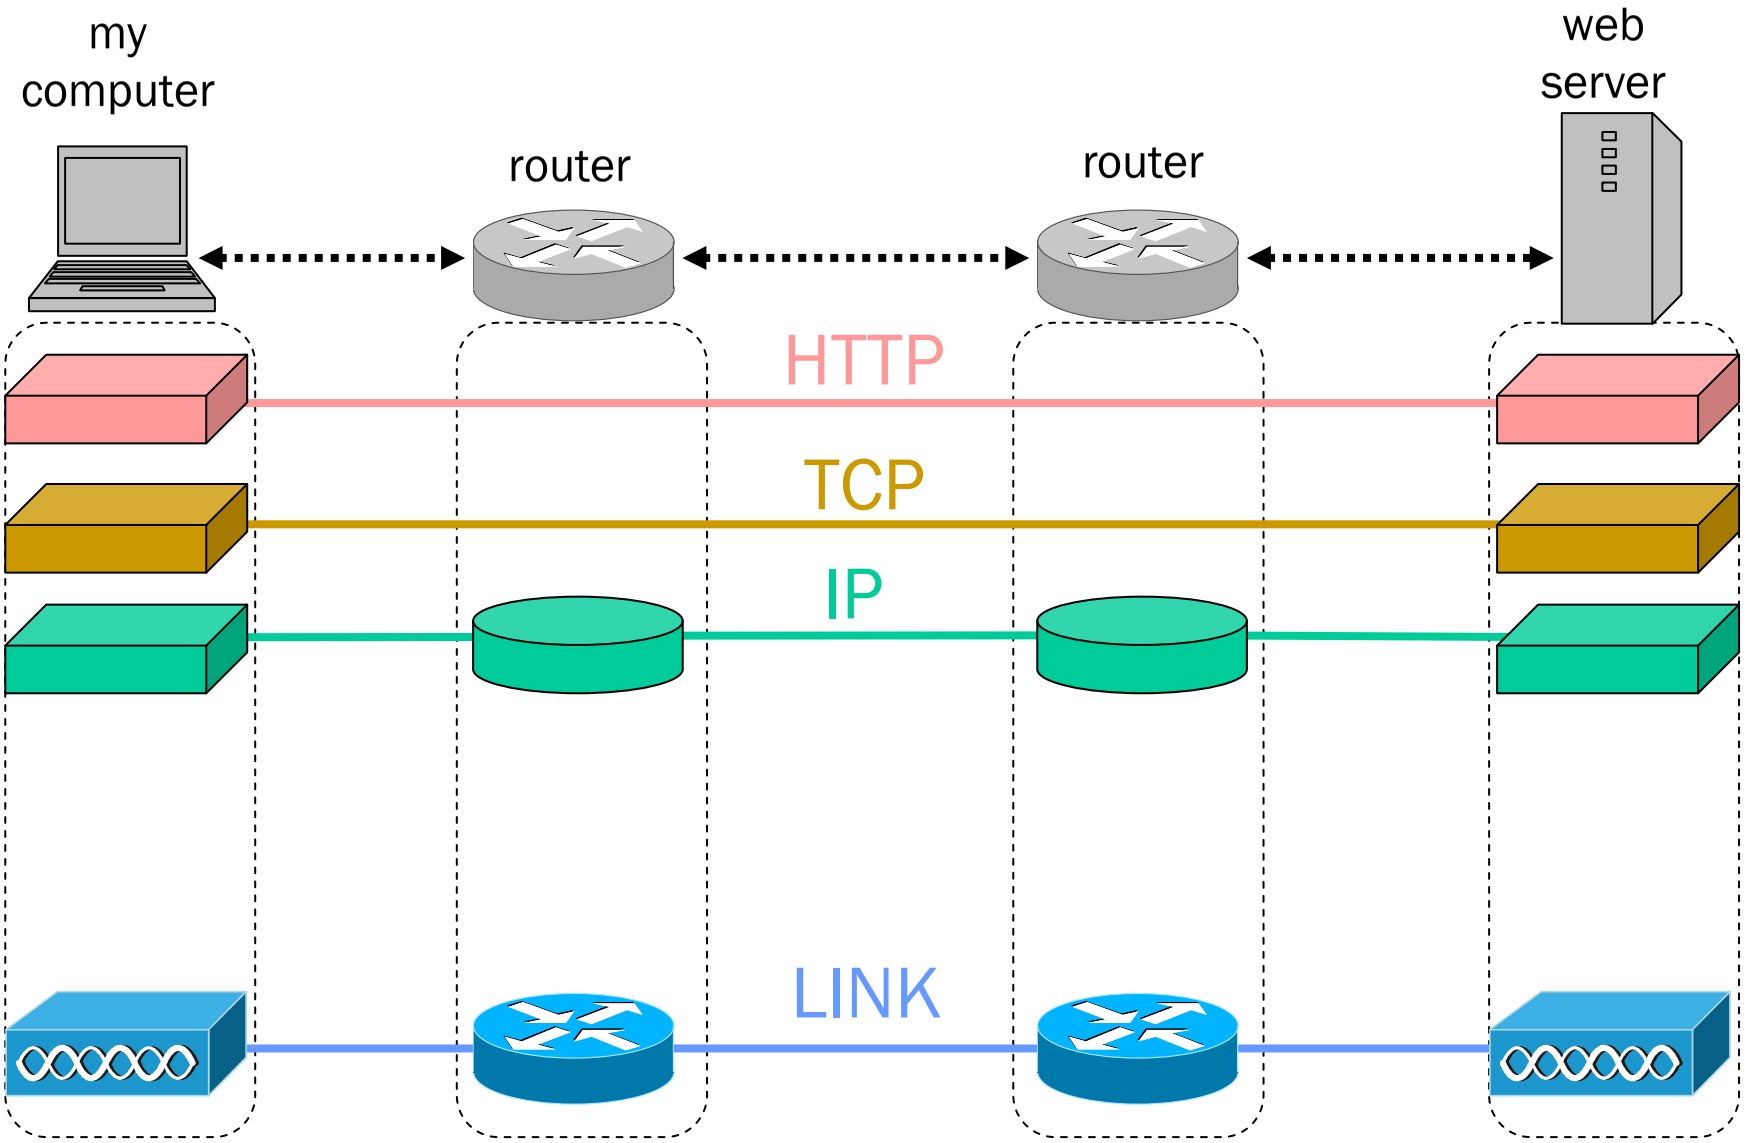

Background image courtesy JHU, applied physics labs

http://www.cs.washington.edu/research/networking/rocketfuel/

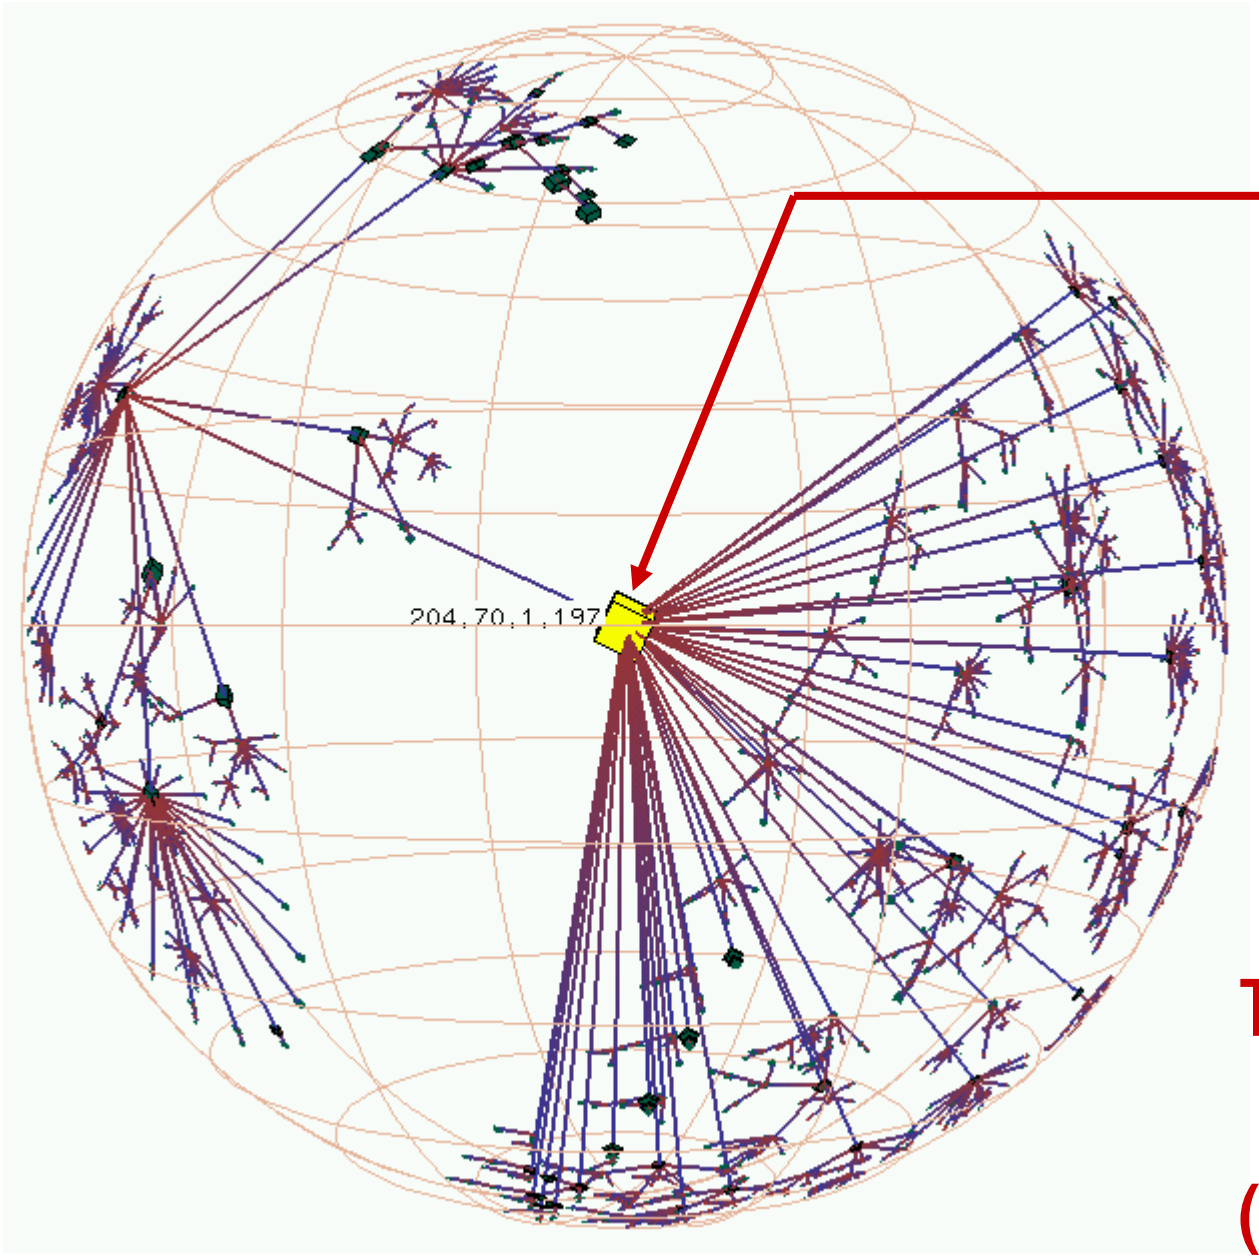

- www.savvis.net
- managed IP and hosting company
- founded 1995
- offering “private IP with ATM at core”

This “node” is an entire network! (not just a router)

http://www.caida.org/tools/measurement/skitter/

HOWEVER: Problems with existing measurements

- traceroute-based measurements are ambiguous
  - traceroute is strictly about IP-level connectivity
  - traceroute cannot distinguish between high connectivity nodes that are for real and that are fake and due to underlying Layer 2 (e.g., Ethernet, ATM) or Layer 2.5 technologies (e.g., MPLS)
- traceroute-based measurements are inaccurate
  - Requires some guesswork in deciding which IP addresses/interface cards refer to the same router ("alias resolution" problem)
- traceroute-based measurements are incomplete/biased
  - IP-level connectivity is more easily/accurately inferred the closer the routers are to the traceroute source(s)
  - Node degree distribution is inferred to be of the power-law type even when the actual distribution is not

Claims about the Internet router-level topology

- How to lie with statistics …
  - Scale-free (power-law node degree distribution)
- (White) lies …
  - Preferential attachment-type models
- Damned lies …
  - Achilles' heel
  - Fragile/vulnerable to targeted node removal

How to lie with statistics …

Given: Samples from an exponential distribution

Want: Claim power law behavior

Recipe: Use size-frequency plots!

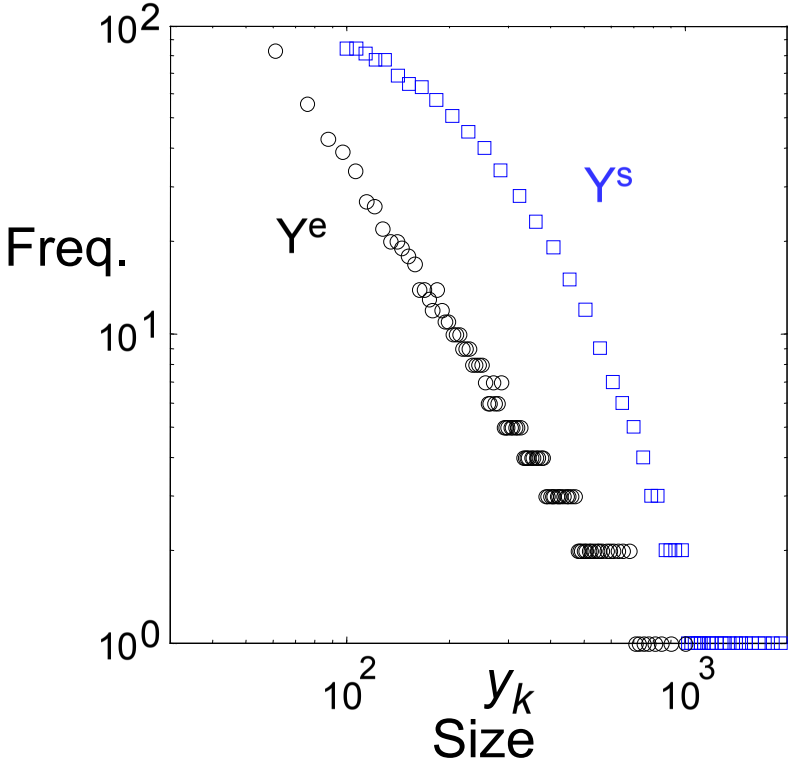

Given: Samples from a Pareto distribution with α=1.0

Want: Claim power law with α=1.5

Recipe: Use size-frequency plots!

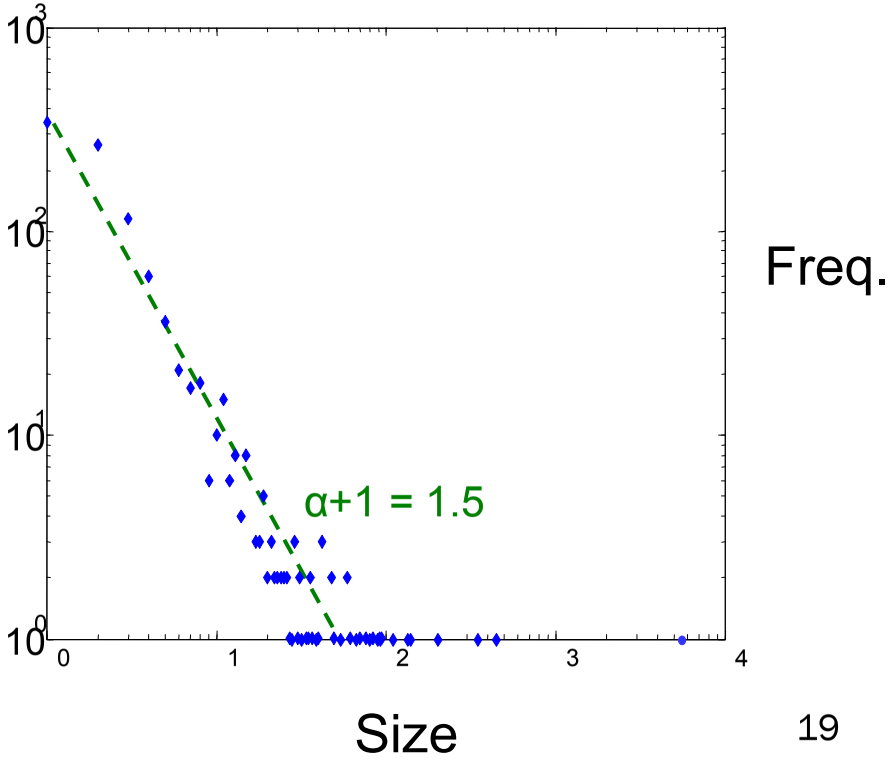

Size-Frequency vs. Size-Rank Plots
or
Non-cumulative vs. Cumulative

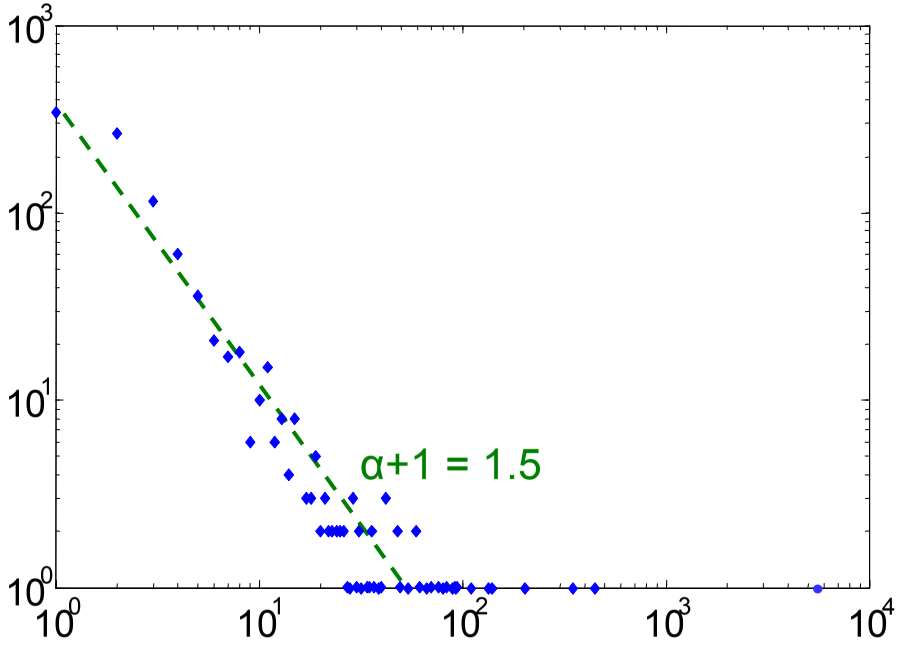

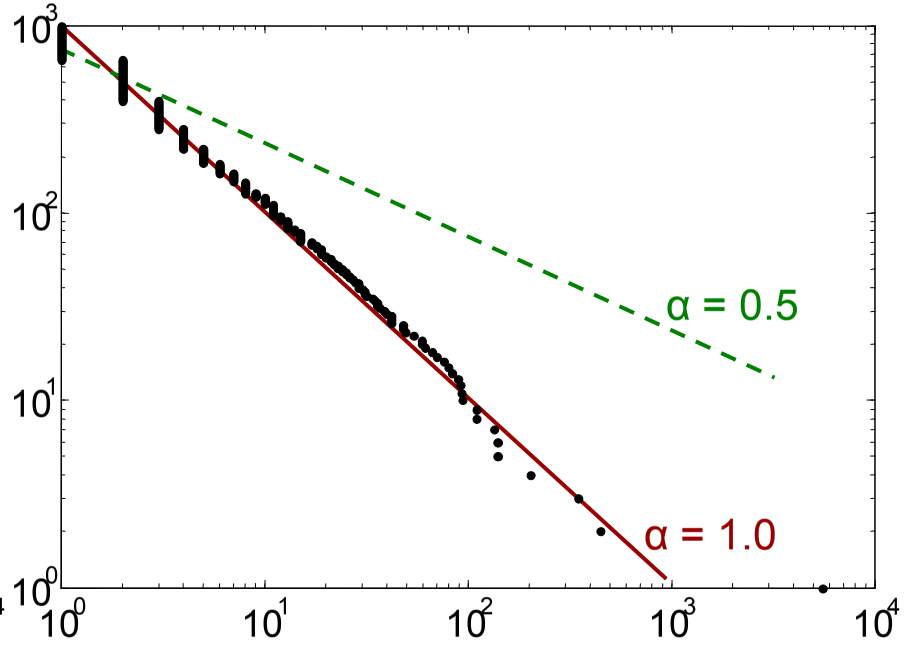

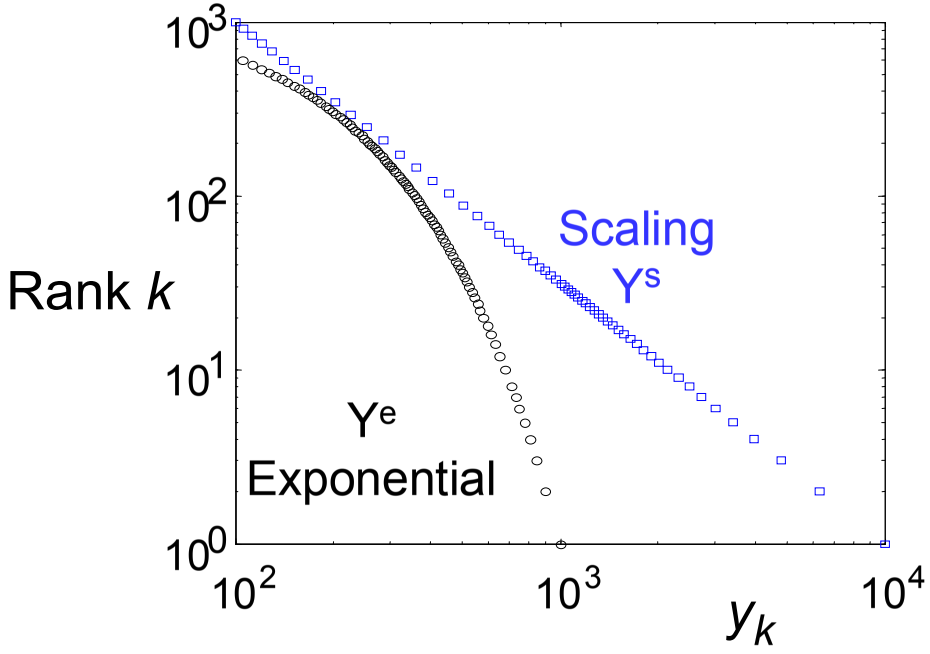
Rank k
10^3
10^2
10^1
10^0
Scaling
Y^s
Y^e
Exponential
10^2
10^3
10^4
y_k

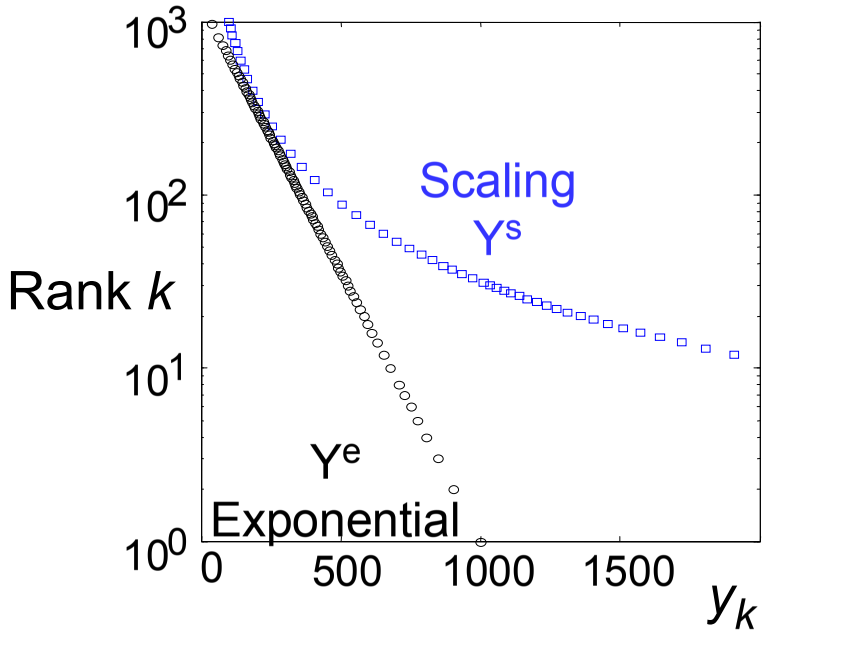
Rank k
10^3
10^2
10^1
10^0
Scaling
Y^s
Y^e
Exponential
0
500
1000
1500
y_k

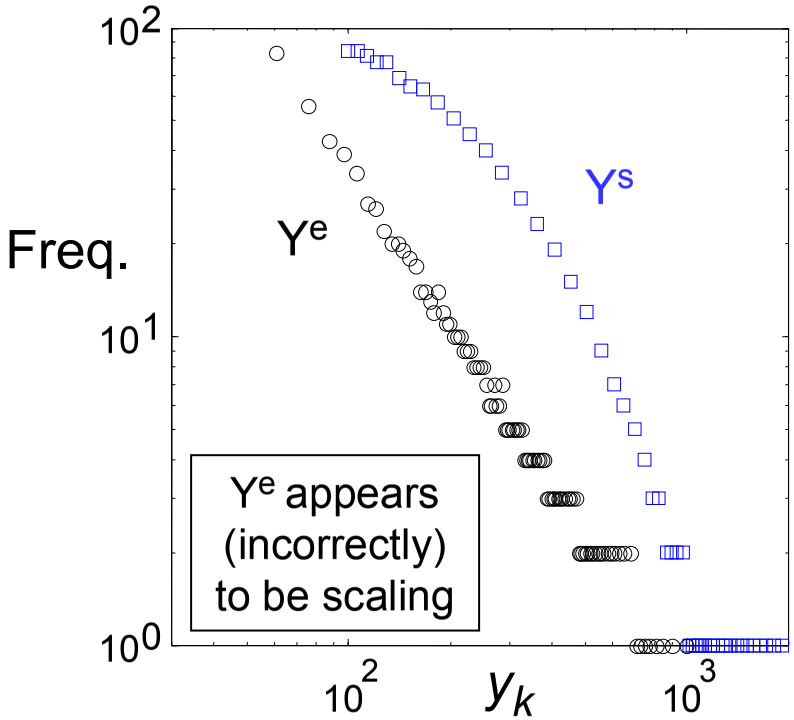
Freq.
10^2
10^1
10^0
Y^s
Y^e
Y^e appears (incorrectly) to be scaling
10^2
y_k
10^3

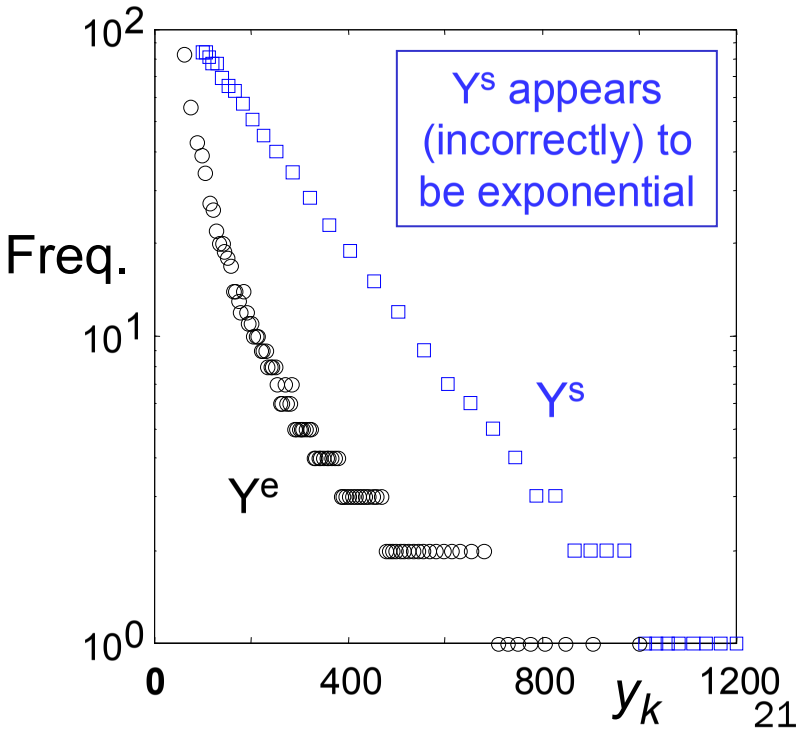
Freq.
10^2
10^1
10^0
Y^s appears (incorrectly) to be exponential
Y^s
Y^e
0
400
800
1200
y_k

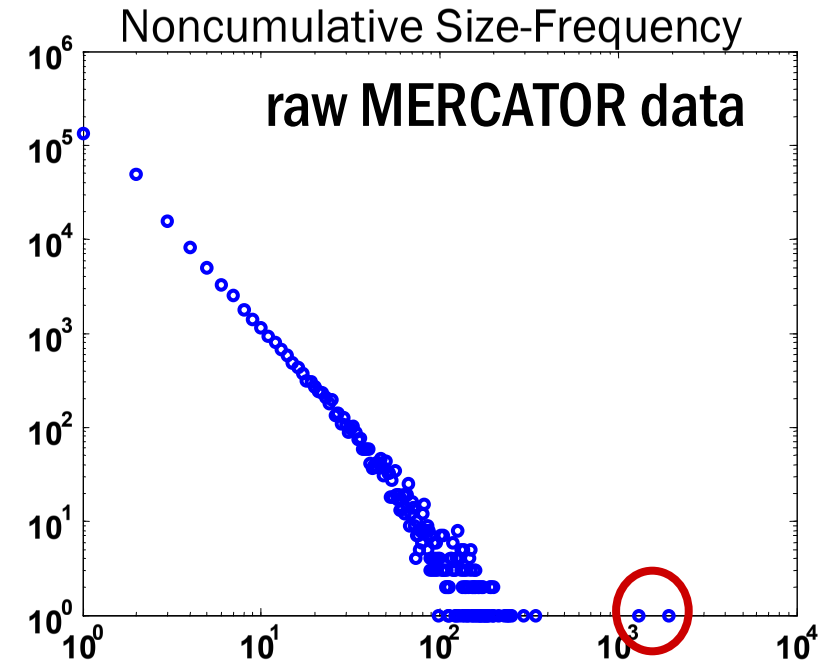
Noncumulative Size-Frequency
raw MERCATOR data

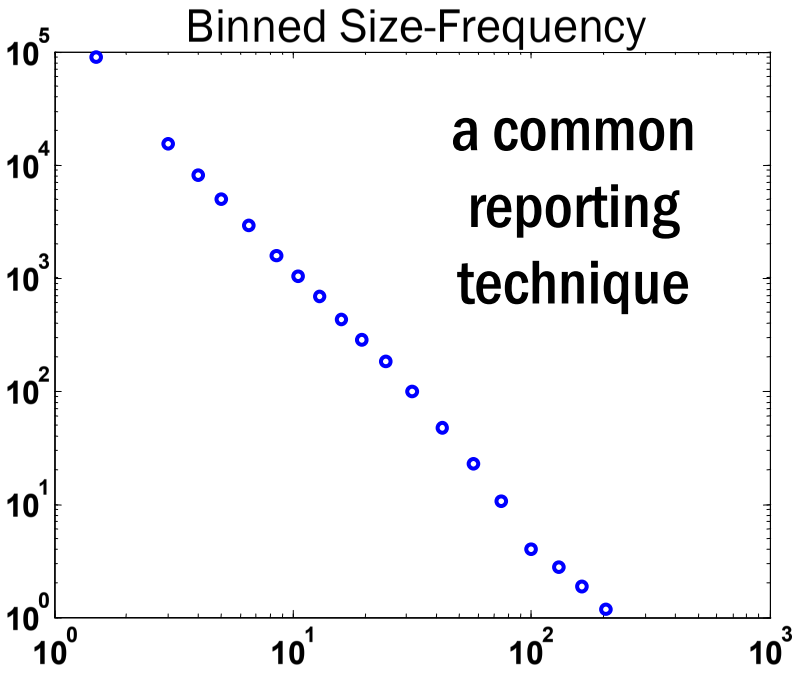
Binned Size-Frequency
a common reporting technique

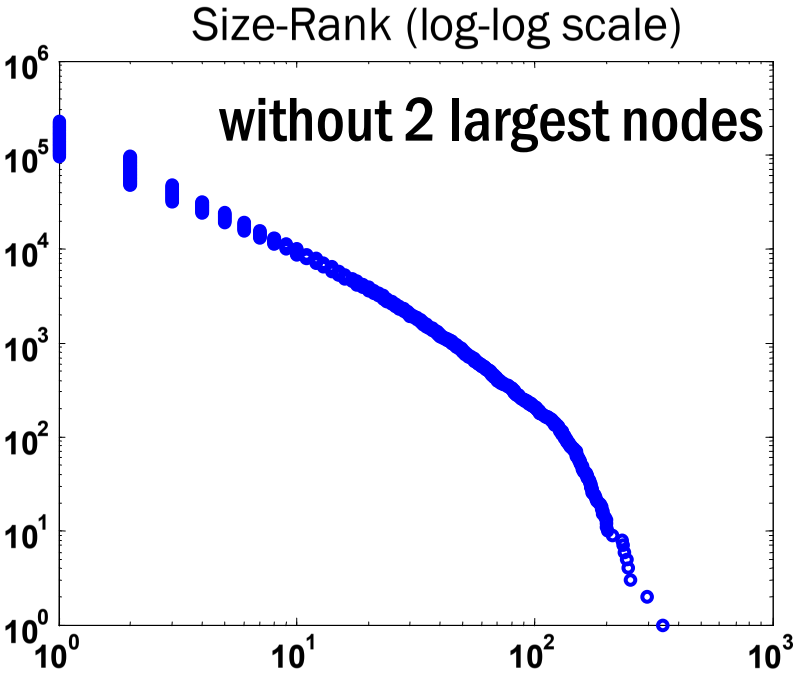
Size-Rank (log-log scale)
without 2 largest nodes

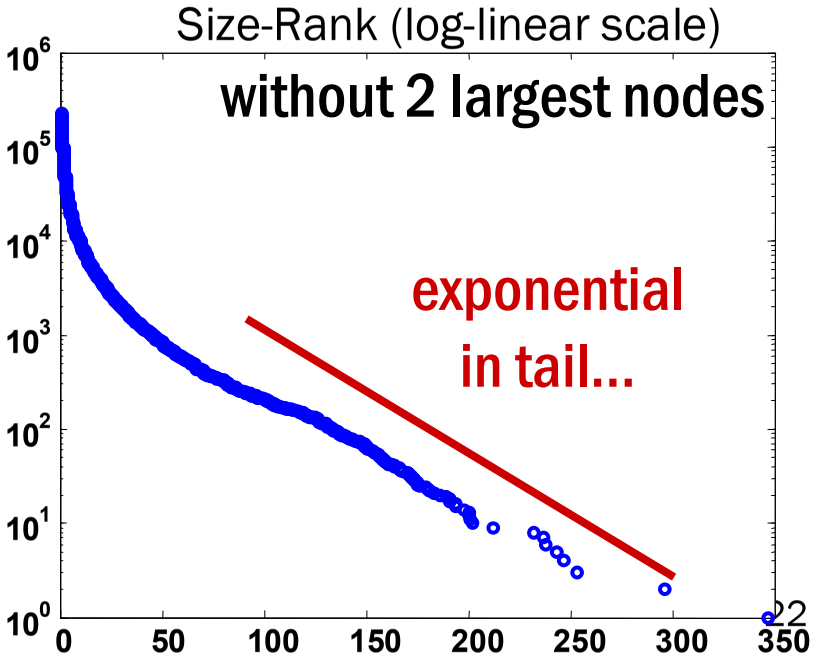
Size-Rank (log-linear scale)
without 2 largest nodes
exponential in tail...

How to avoid such fallacies: Taking Model validation serious …

- Fact 1
  - For one and the same observed phenomenon, there are usually many different explanations/models
  - All models are wrong, but some are “damned lies”
  - The ability to reproduce a few graph statistics does not constitute “serious” model validation
- Fact 2
  - Recent alternatives to PA-type models: HOT-models
  - Key features of HOT models
    - Consistent with existing ISP router-level topologies
    - Consistent with existing technologies
    - Consistent with (complementary) measurements
    - Node degree distribution is a non-issue

Observation#1: Not consistent with real-world network topologies

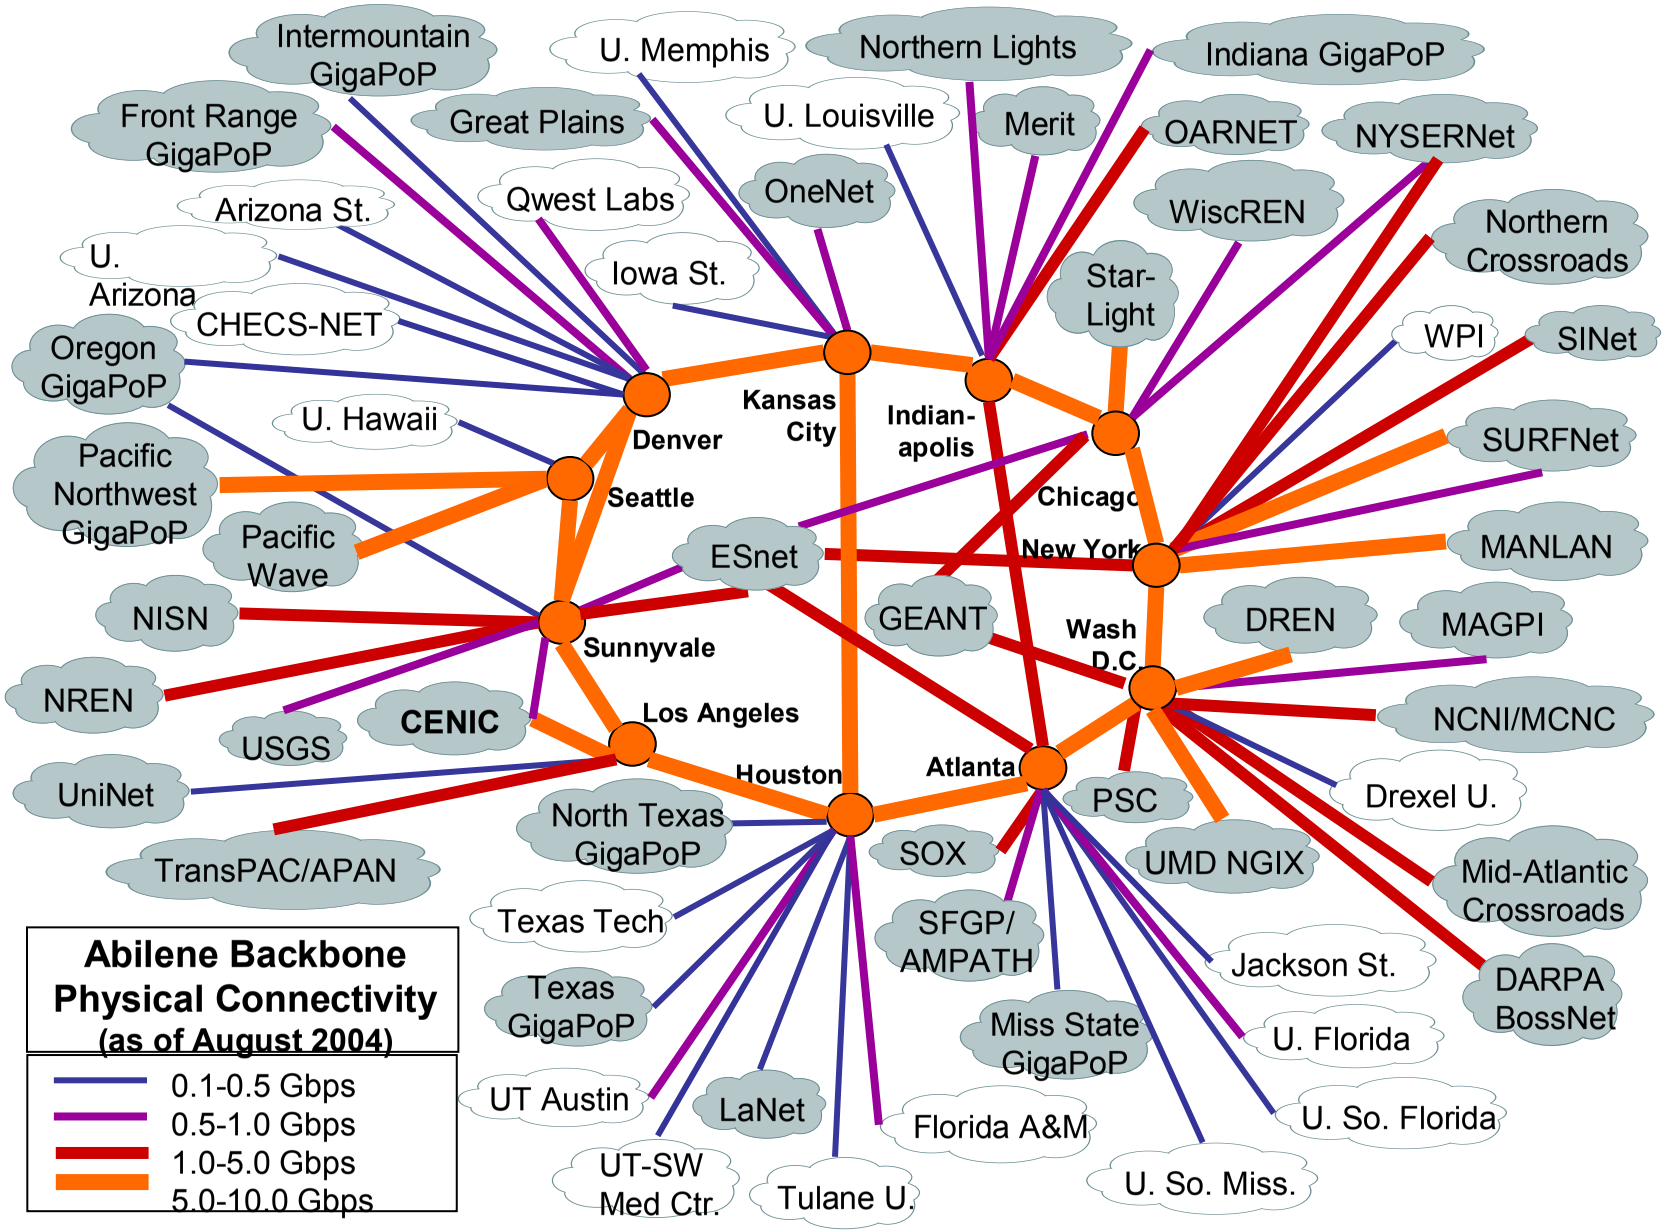

Cisco 12000 Series Routers

- Modular in design, creating flexibility in configuration.
- Router capacity is constrained by the number and speed of line cards inserted in each slot.

Chassis	Rack size	Slots	Switching Capacity
12416	Full	16	320 Gbps
12410	1/2	10	200 Gbps
12406	1/4	6	120 Gbps
12404	1/8	4	80 Gbps

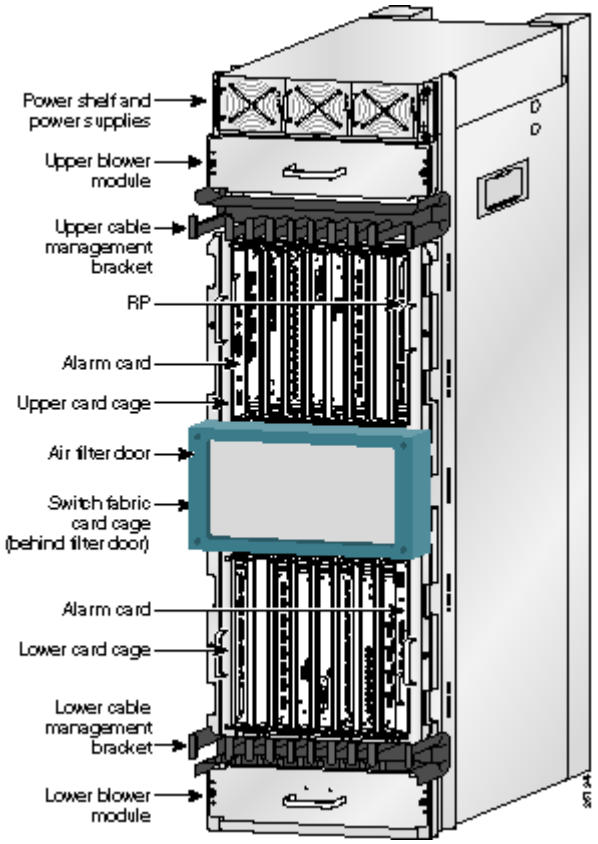

Source: www.cisco.com

Router Technology Constraint

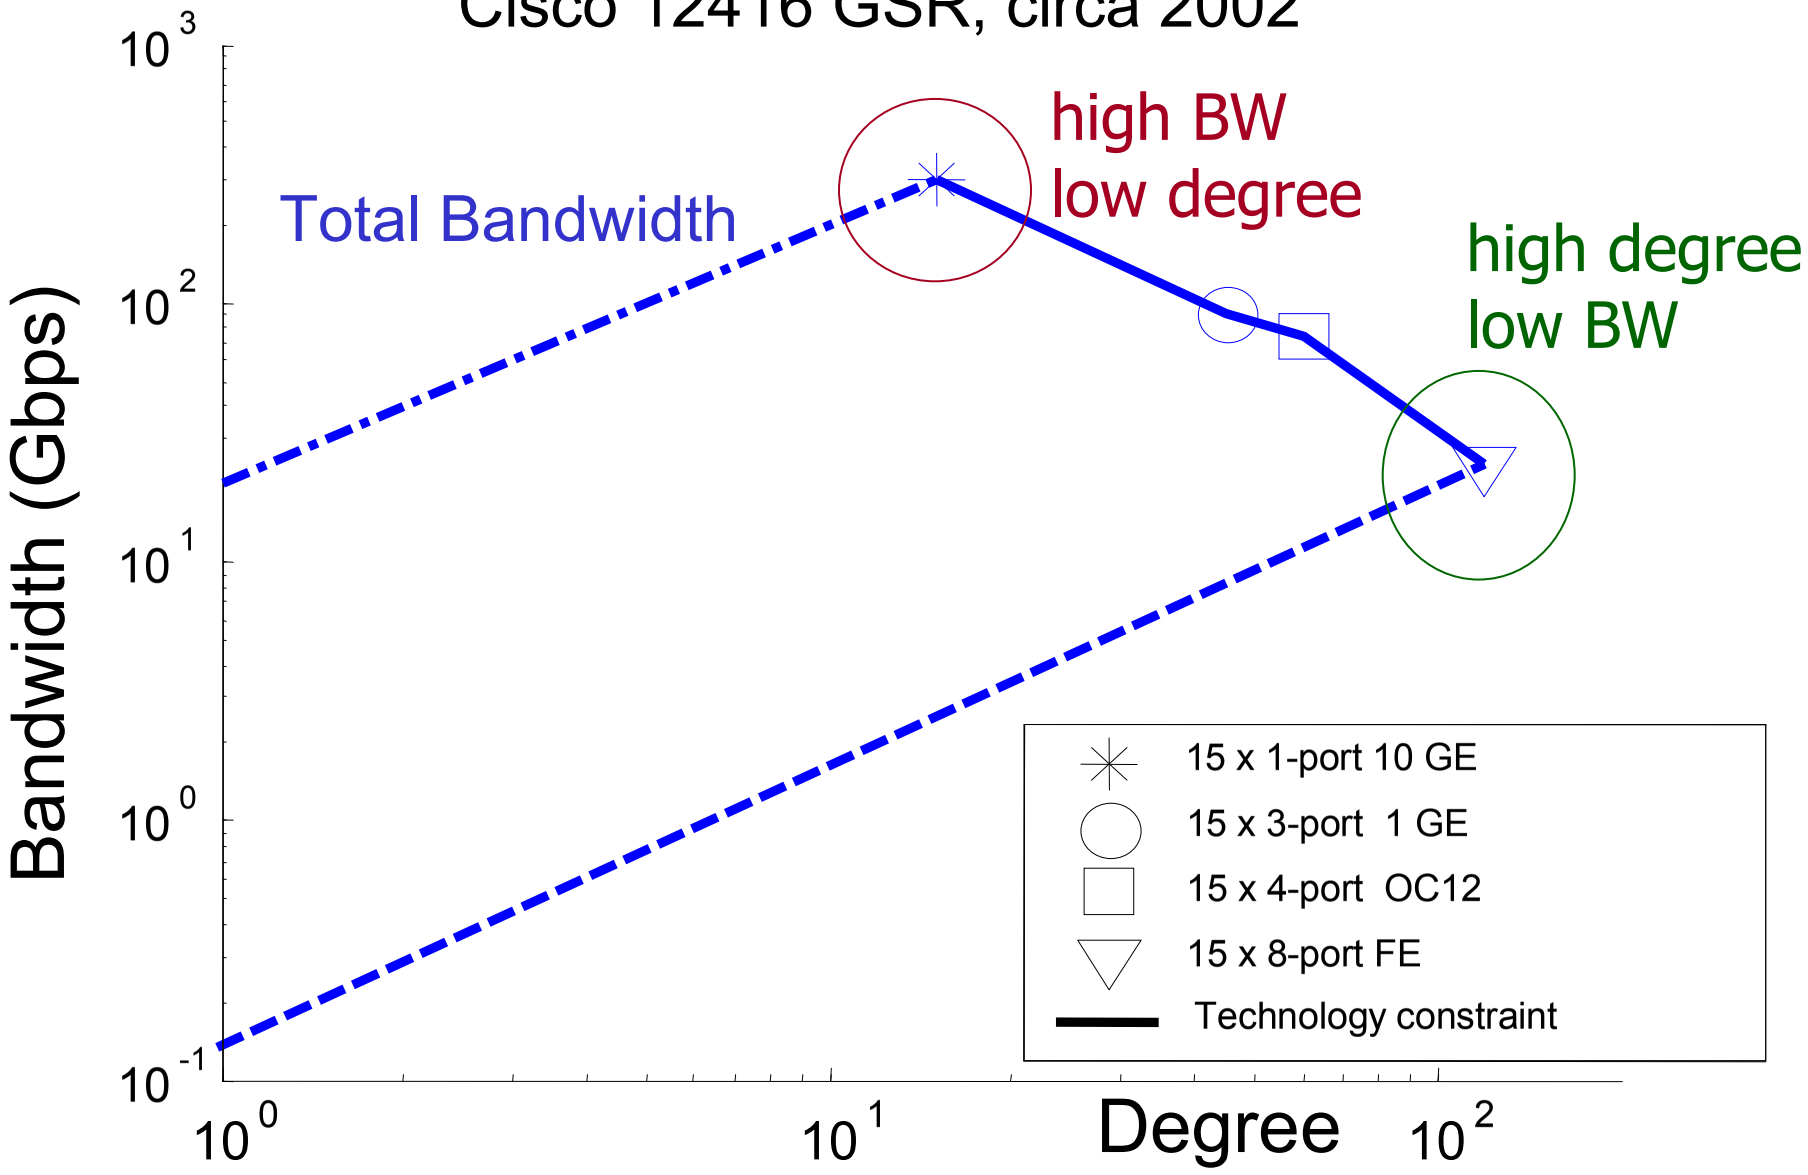

Router Deployment: Abilene and CENIC

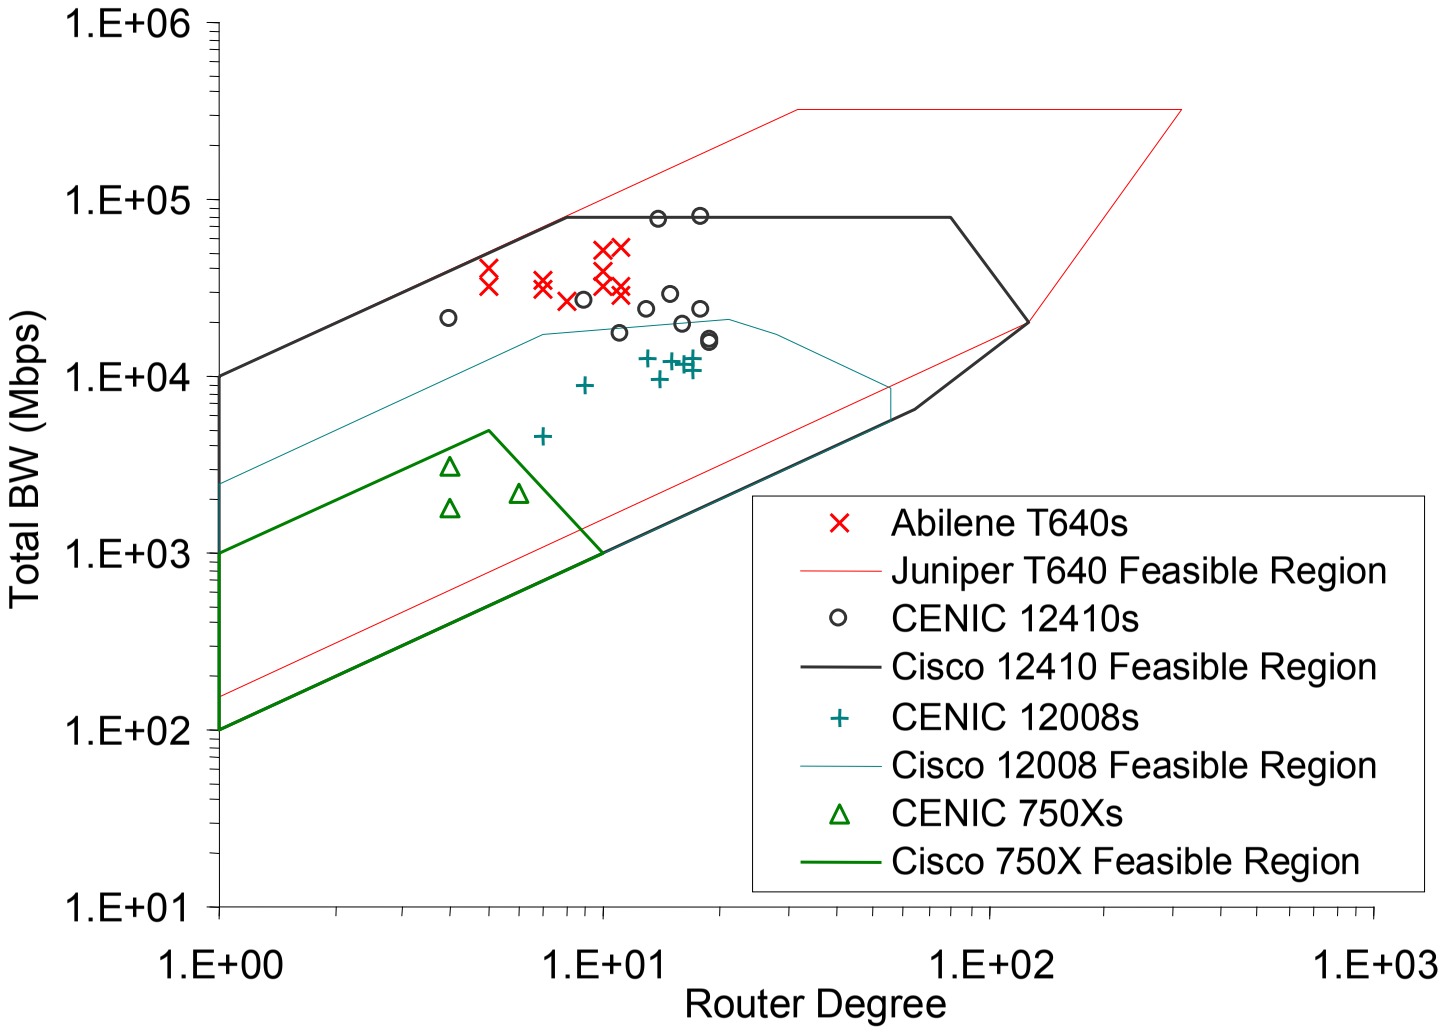

AT&T Router Deployment (c.2003)

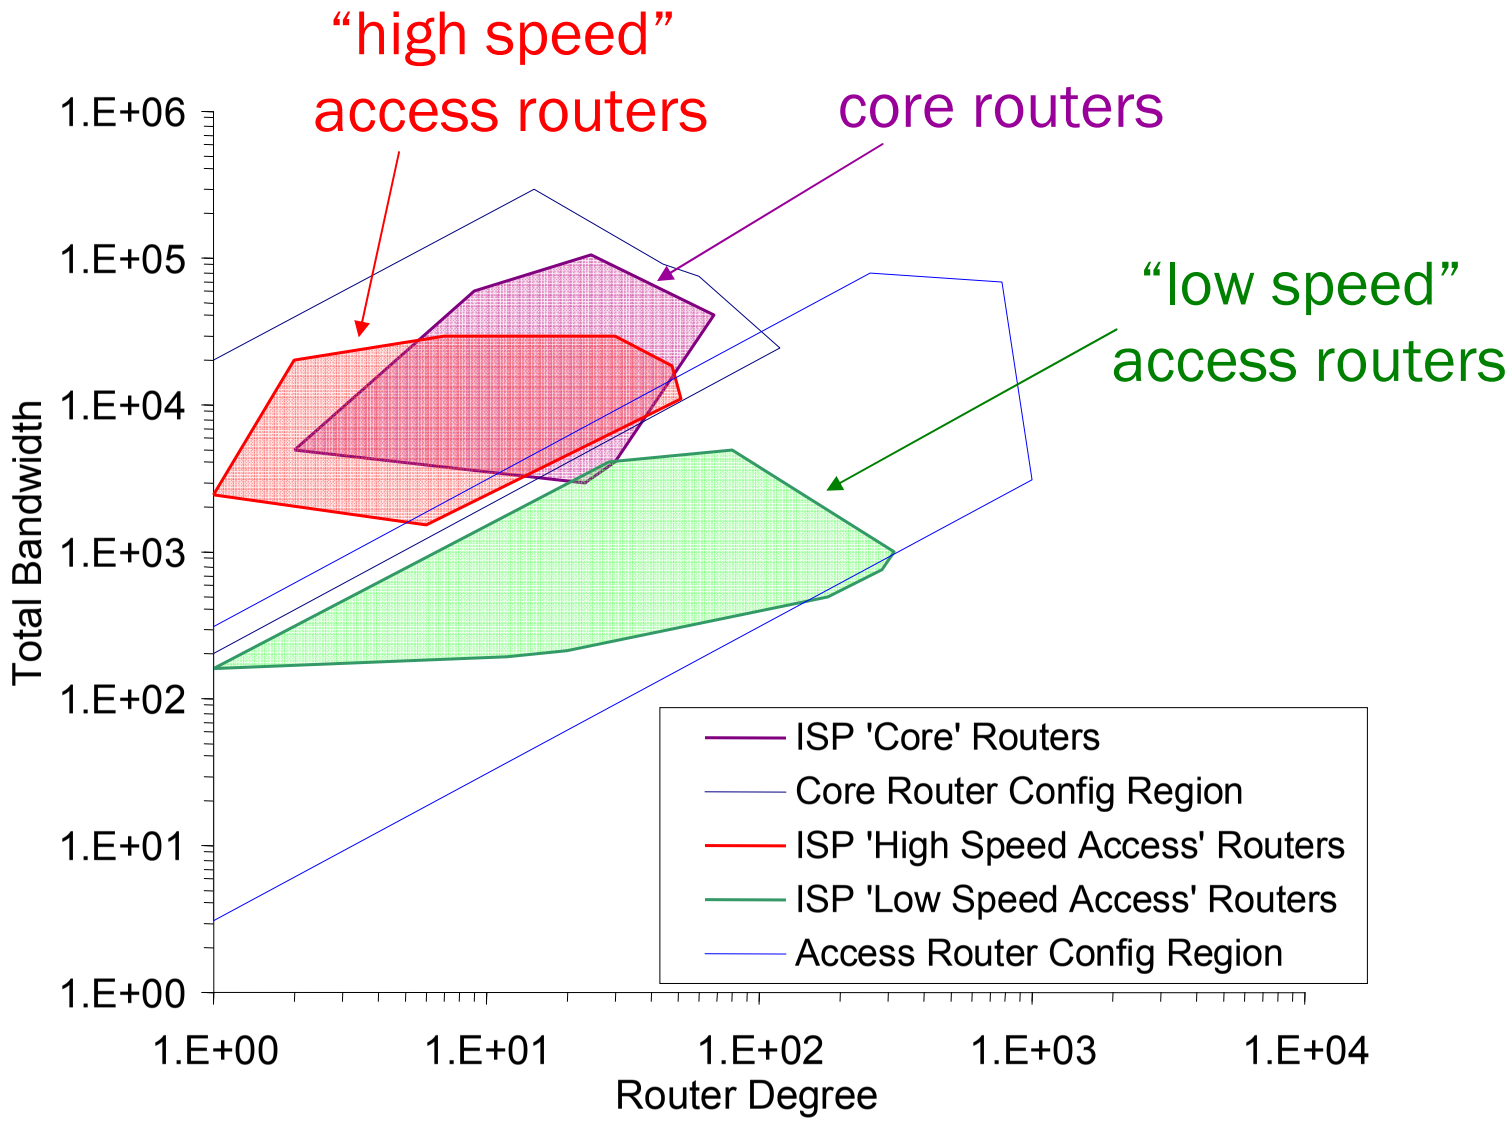

Litmus Test for proposed Network Models

- Make node degree distribution a non-issue
  - Good reasons
    - High-quality data but low variability (e.g., exponential)
    - Low-quality data
    - High-quality data and high variability (e.g., power-laws)
  - PA-type models
    - dead on arrival
  - Only reasonable alternative
    - Bring in and rely on domain knowledge
- What new kinds of measurements does the proposed model suggest for the purpose of model validation
  - PA-type models: none
  - HOT models: get data on existing router technology

On Measuring the Internet's AS-level Topology

AS-Level Topology

- Nodes = (sets of) entire networks (Autonomous Systems or ASes)
- Links = peering relationships between ASes
- Really a map of economic or business relationships, not of physical connectivity

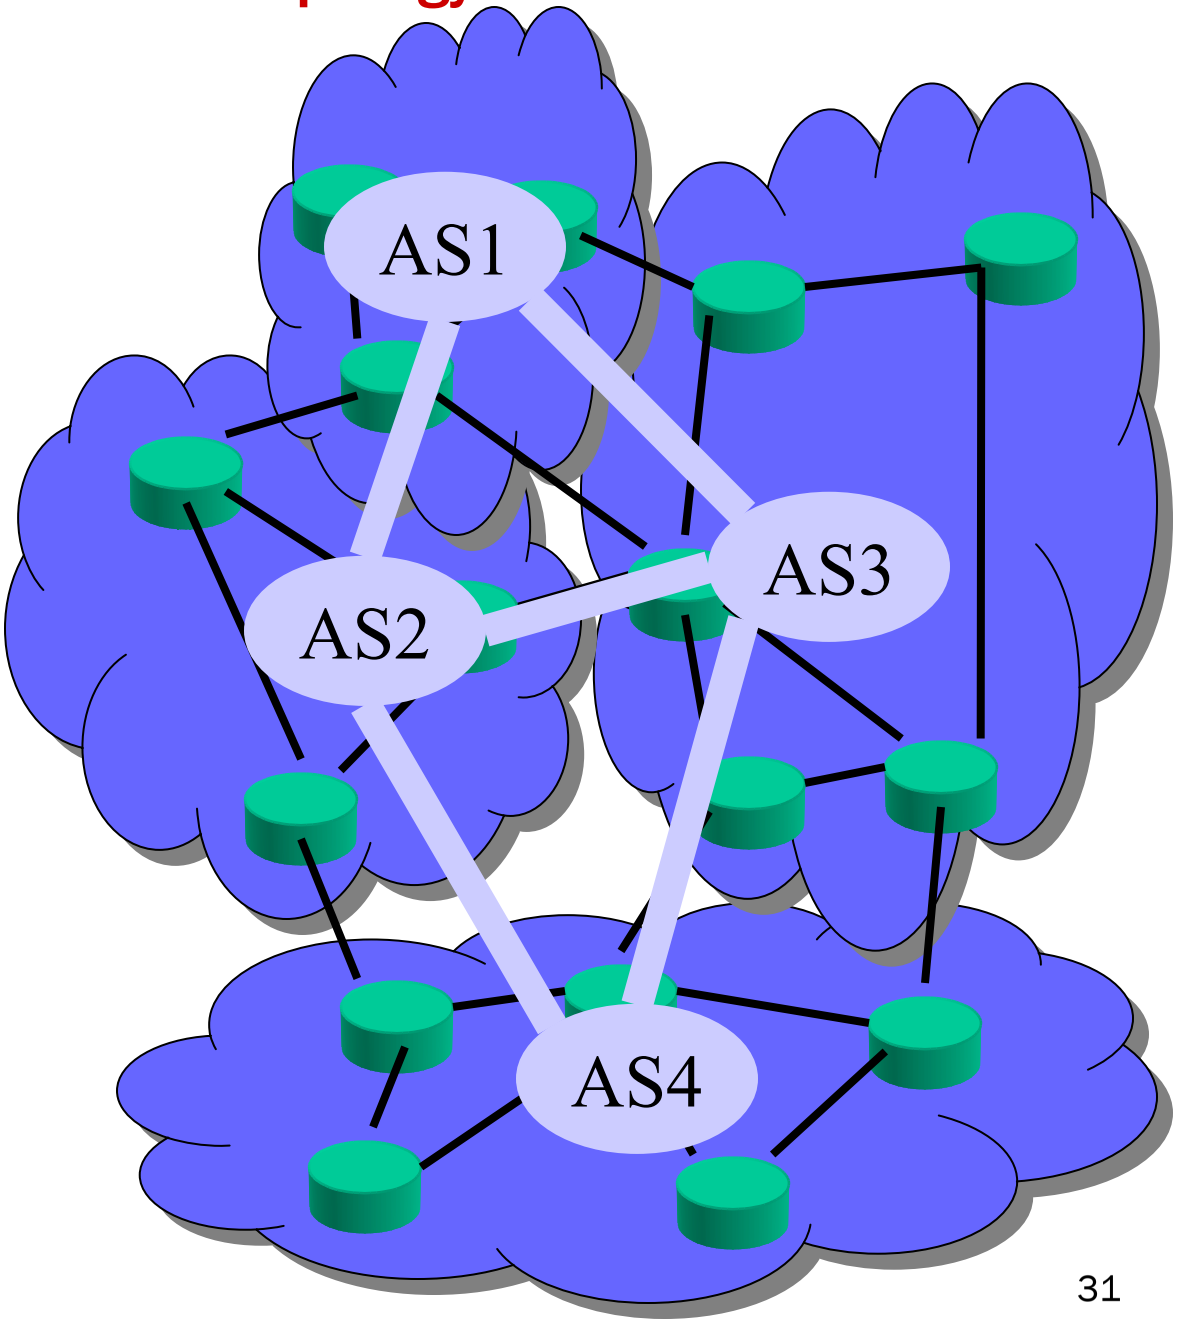

(Annotated) AS-Level Topology

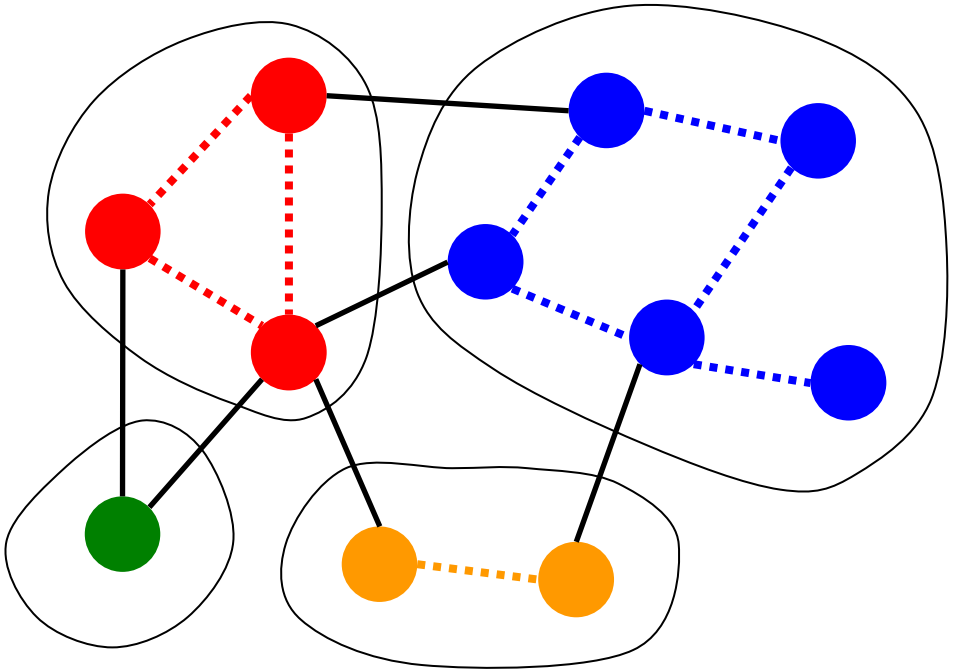

Router
Intra-domain router-level link
Inter-domain router-level link

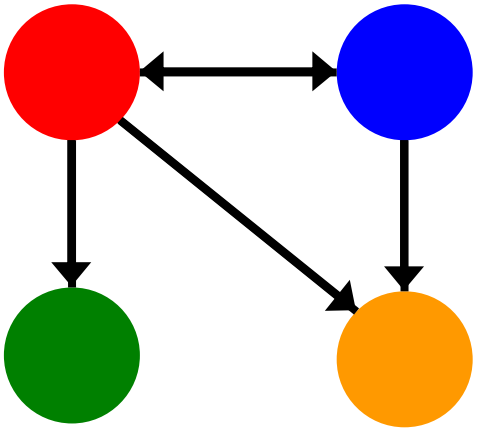

Autonomous System (AS)
Peering relationship
Provider-customer type
Peer-to-peer type

Note 1: AS Graphs Obscure Physical Connectivity!

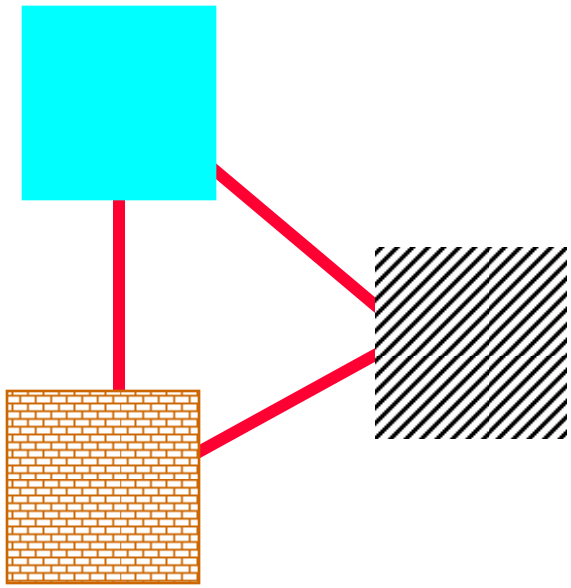

The AS graph may look like this.

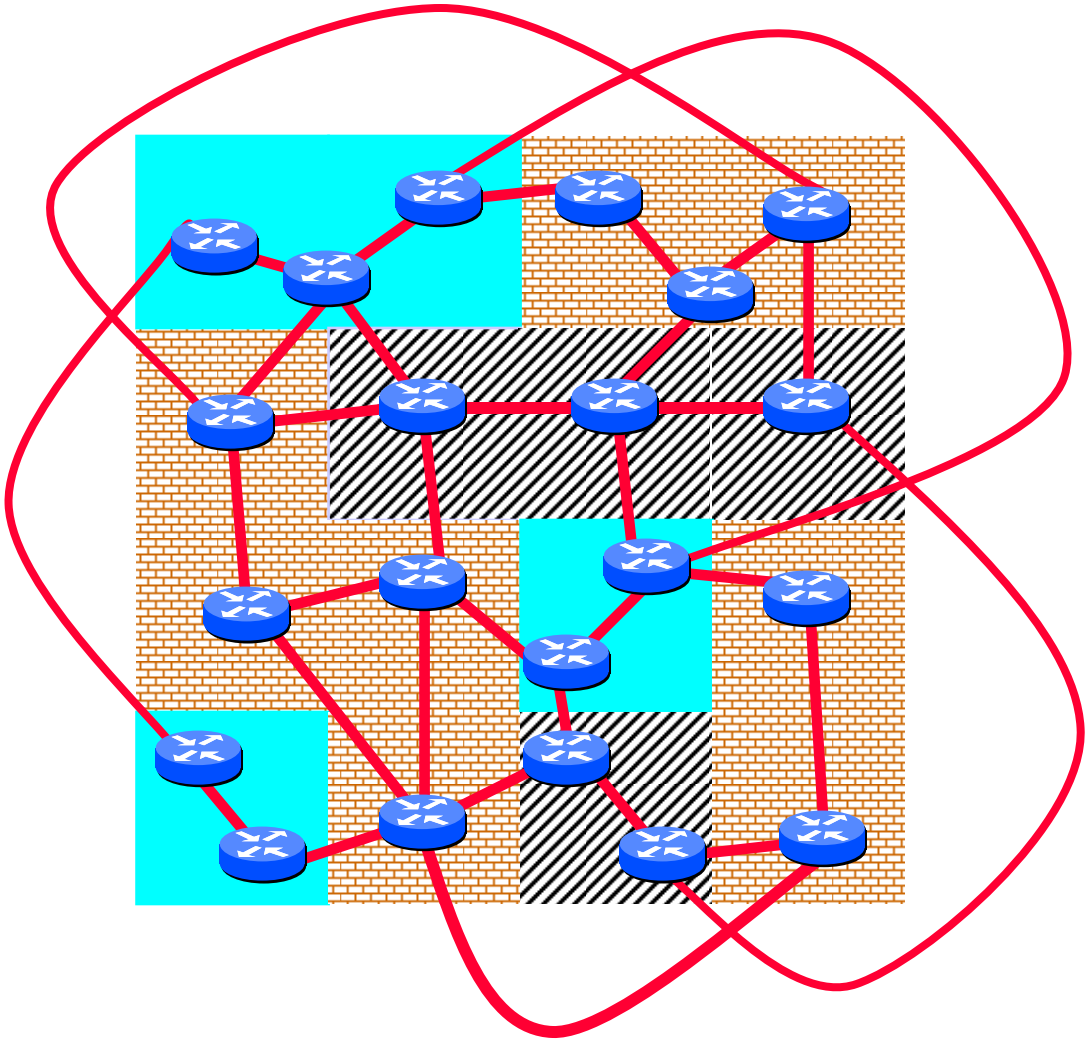

Reality may be closer to this…

Courtesy Tim Griffin

Note 2: AS Graphs obscure Geography!

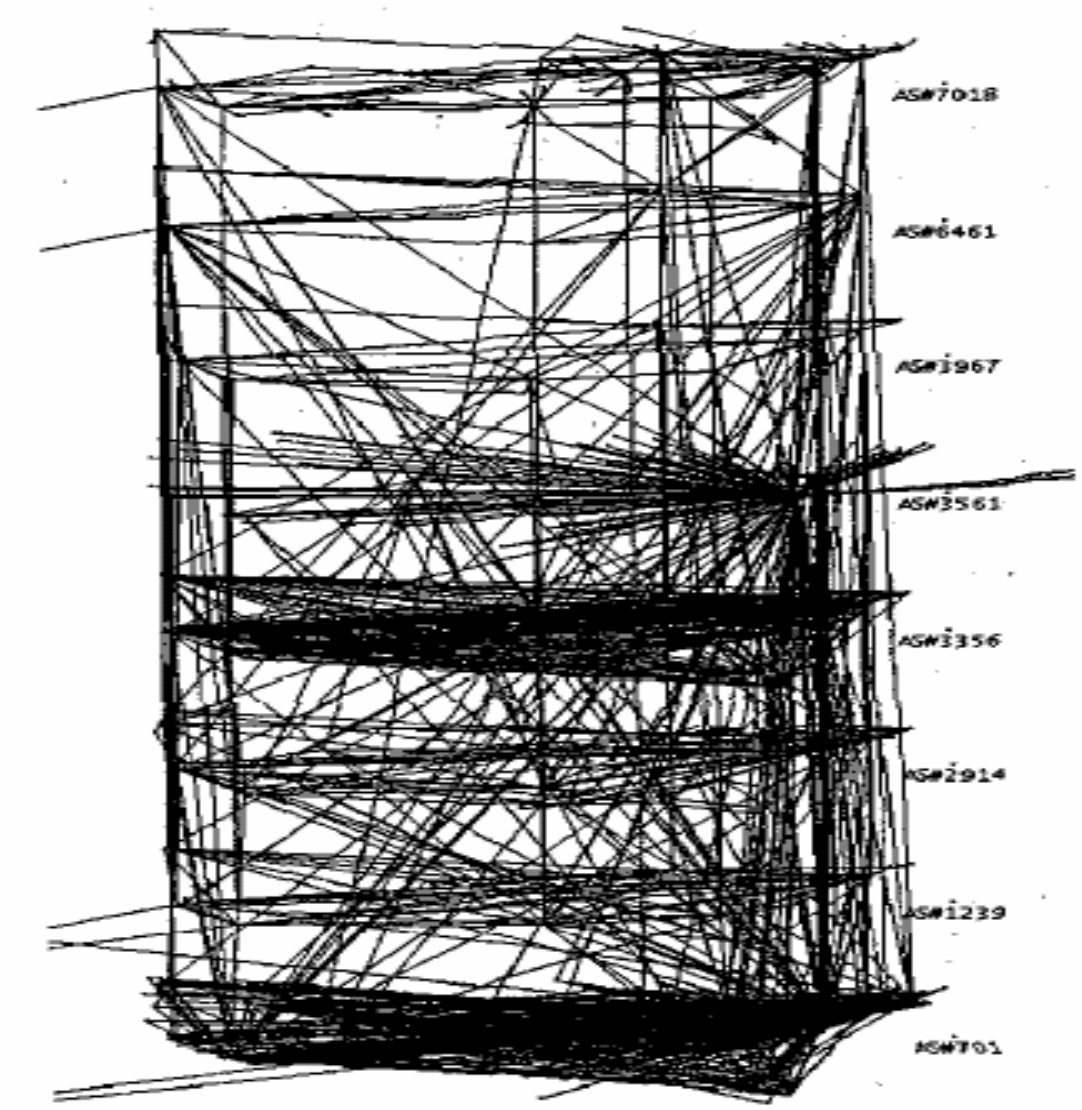

Liljenstam, Liu, and Nicol (2003)

On Measuring the Internet's AS-level Topology

- BGP routing tables/updates
  - RouteViews (Univ. of Oregon)
  - RIPE (Europe)
  - E.g., 129.223.224.0/19 7018 701 4637 1221
- Traceroute measurements
  - Skitter (CAIDA, US)
  - DIMES (EU-project)
- Other available sources
  - Public databases (WHOIS)
  - Looking glass sites, additional routing tables

HOWEVER: Problems with existing measurements

- BGP-based measurements
  - Incomplete: up to 40-50% of existing links are missing
  - Ambiguous: requires significant guesswork in deciding whether a "new" node or link is genuine
  - Inaccurate: heuristics for inferring peering relationships
- WHOIS database
  - Unreliable: ISP's have no incentives to register
  - Incomplete: ISP's have no incentives to provide info
  - Out-of-date: stale data
- Traceroute-based measurements
  - See earlier

Conventional Approaches to AS-level Topology Modeling

- Step 1: Take the available measurements at face value
- Step 2: Analyze the data as if they could provide the ground truth about the Internet's actual AS-level connectivity structure
- Step 3: Propose a random graph model or construction that describes/fits the inferred AS maps well
- Step 4: Argue for the validity of the proposed model on the basis that it is capable of reproducing certain empirically observed properties of the inferred AS maps

Criticism of conventional modeling approach

- Measurements
  - Connectivity-related Internet measurements are of limited quality
  - BGP is not a mechanism by which ASs distribute connectivity information, but is a protocol by which ASs distribute the reachability of their networks via a set of routing paths that have been chosen by other ASs in accordance with their policies.
- Modeling
  - Inferred AS maps are in general dubious or useless, unless they are accompanied by strong robustness results that state whether or not the observed properties are insensitive to the known ambiguities inherent in the underlying measurements.
  - Chang et al. (2004)
- Model validation
  - Which of the observed properties does a proposed model have to satisfy before it is deemed “valid”?
  - The ability to reproduce a few graph statistics does not constitute “serious” model validation

On the Completeness of BGP-Inferred Connectivity Maps

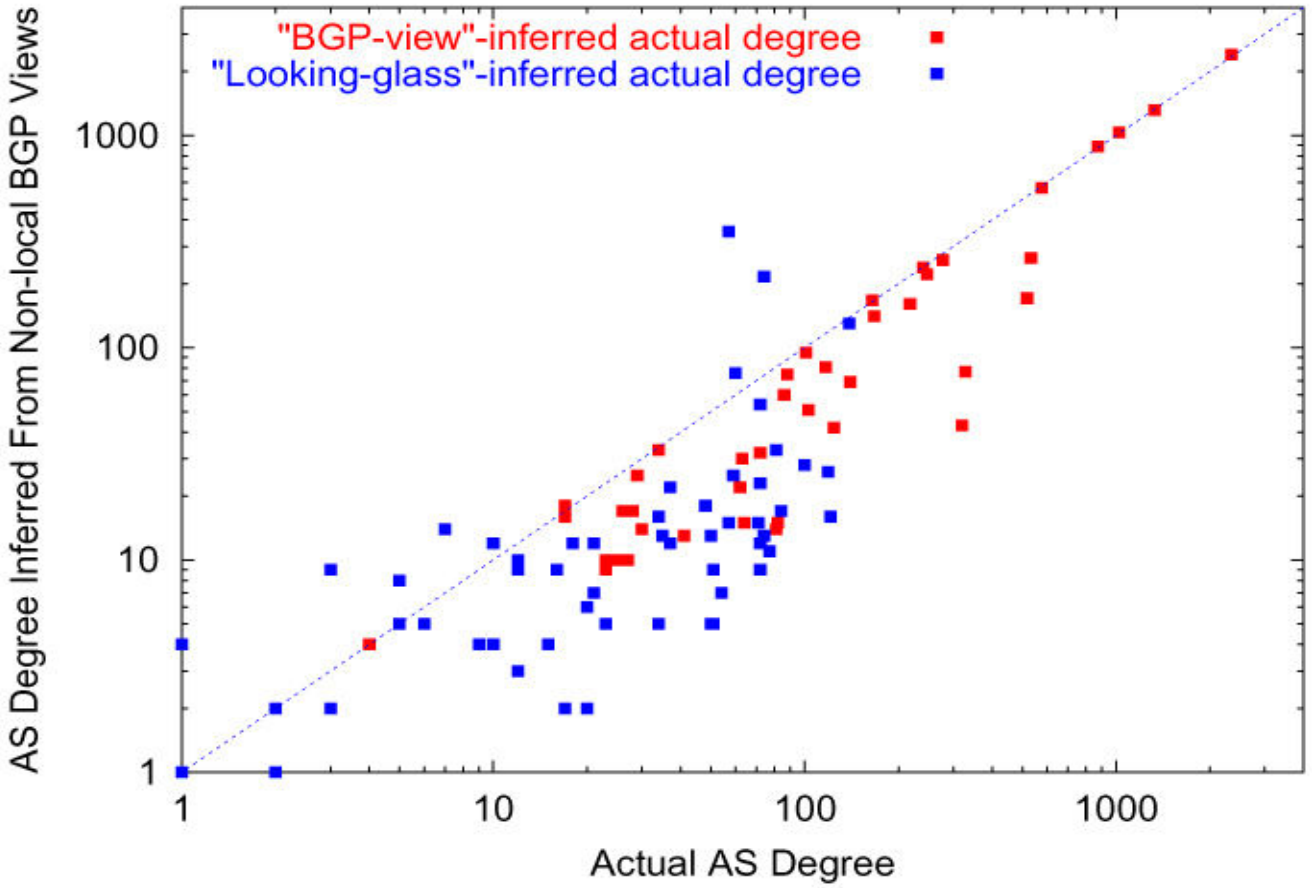

Claims about the Internet AS-level topology

- How to lie with statistics …
  - Scale-free (power-law node degree distribution)
- (White) lies (?)
  - Preferential attachment-type models
- Damned lies …
  - Achilles' heel
  - Fragile/vulnerable to targeted node removal

A New Modeling Paradigm for AS-level Topology Modeling

- Engineering perspective
  - Surely, deciding on whether or not to establish what type of peering relationship and with whom is not the outcome of a series of chance experiments conducted by the different ASs, but is largely based on economic arguments.
  - AS relationships are meaningless without the notion of traffic
- Engineering approach
  - Node degree distribution is a non-issue
  - Optimization of tradeoffs between multiple functional objectives of networks
  - Subject to constraints on their components
  - With an explicit source of uncertainty against which solutions must be tolerant or robust
- Recent alternatives to PA-type models
  - Chang et al. (2006)

Some Implications of this Engineering Perspective

- Dynamics of graphs
  - Evolution of connectivity structures
  - Evolution of (internal) node/link structure
- Dynamics over graphs
  - Traffic dynamics (bytes, packets, flows, …)
- Challenging feedback problem
  - Traffic dynamics/routing impacts network structure
  - Network structure impacts traffic dynamics/routing
- Robustness/fragility considerations only make sense in the context of the broader system, i.e., protocol stack
  - Router-level: Inter-AS routing protocol
  - AS-level: Intra-AS routing protocol

On Measuring the Overlay Topology

- P2P networks
  - Structured (e.g., Kad DHT): Central control
  - Unstructured (e.g., Gnutella): Crawler
- World-Wide-Web (WWW)
  - AltaVista crawls (Broder et al,) in 1999
  - Duration is a couple of weeks

HOWEVER: Problems with existing measurements

- High degree of dynamics of overlay networks
  - Connectivity structure changes underneath the crawler
  - Fast vs. slow crawls
- Enormous size of overlay networks
  - Complete crawls take too long
  - Alternative approach: Sampling
- Issues of sampling bias
  - Due to temporal dynamics of nodes (peers)
  - Due to spatial features of overlay topology

A Reminder

- Modeling in the presence of high-quality data
  - *"All models are wrong … but some are useful"* (G.P.E. Box)

- (Past) modeling in the case of highly ambiguous data
  - Taking the data at face value produces "bad models"
  - *"Bad [models] are potentially important: they can be used to stir up public outrage or fear; they can distort our understanding of our world; and they can lead us to make poor policy choices."* (*J. Best*)

- (Future) modeling in the case of highly ambiguous data
  - Take the ambiguities in the data into account
  - *"When exactitude is elusive, it is better to be approximately right than certifiably wrong."* (*B.B. Mandelbrot*)