# Scaling Phenomena in the Internet

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

# Topics Covered

- The Self-Similar Nature of Internet Traffic
  - The “Physics of Complex Systems” view
  - The Networking view
  - JCD’s dream comes true: 1st INFOCOM paper
- High Variability in Internet Topology
  - The “Physics of Complex Systems” view
  - The Networking view
  - JCD’s dream comes true: 1st SIGCOMM paper
- A Look Ahead

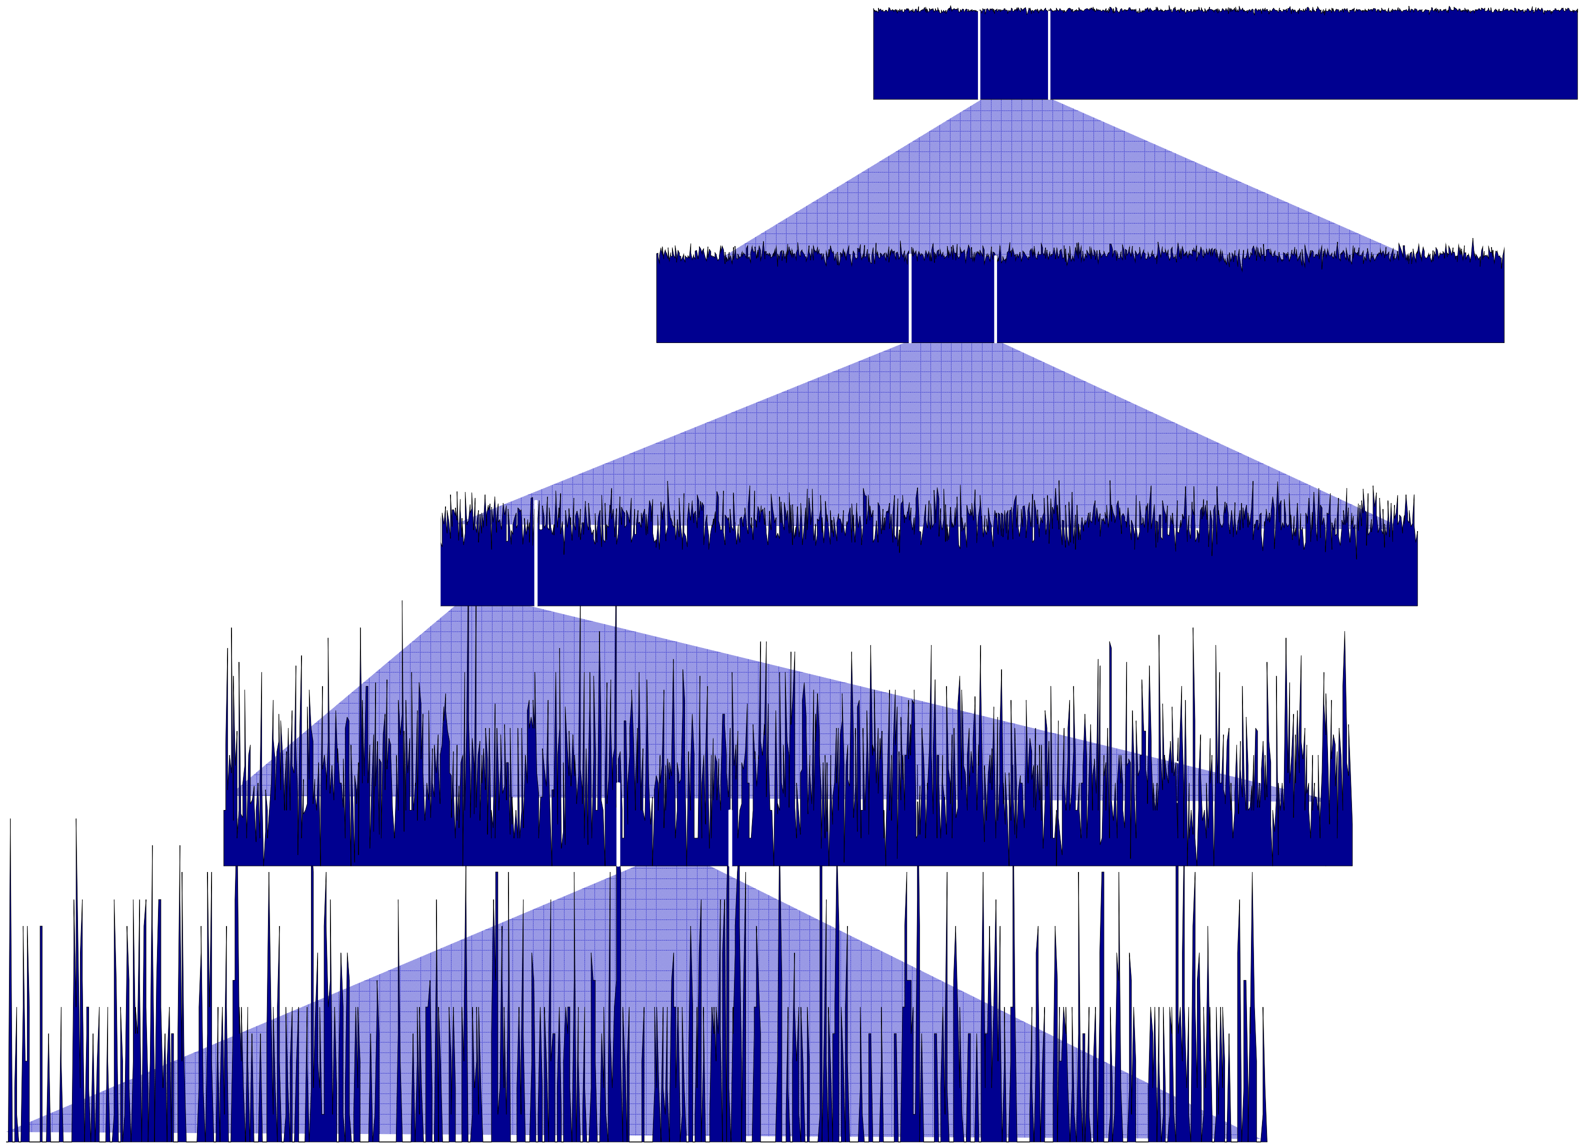

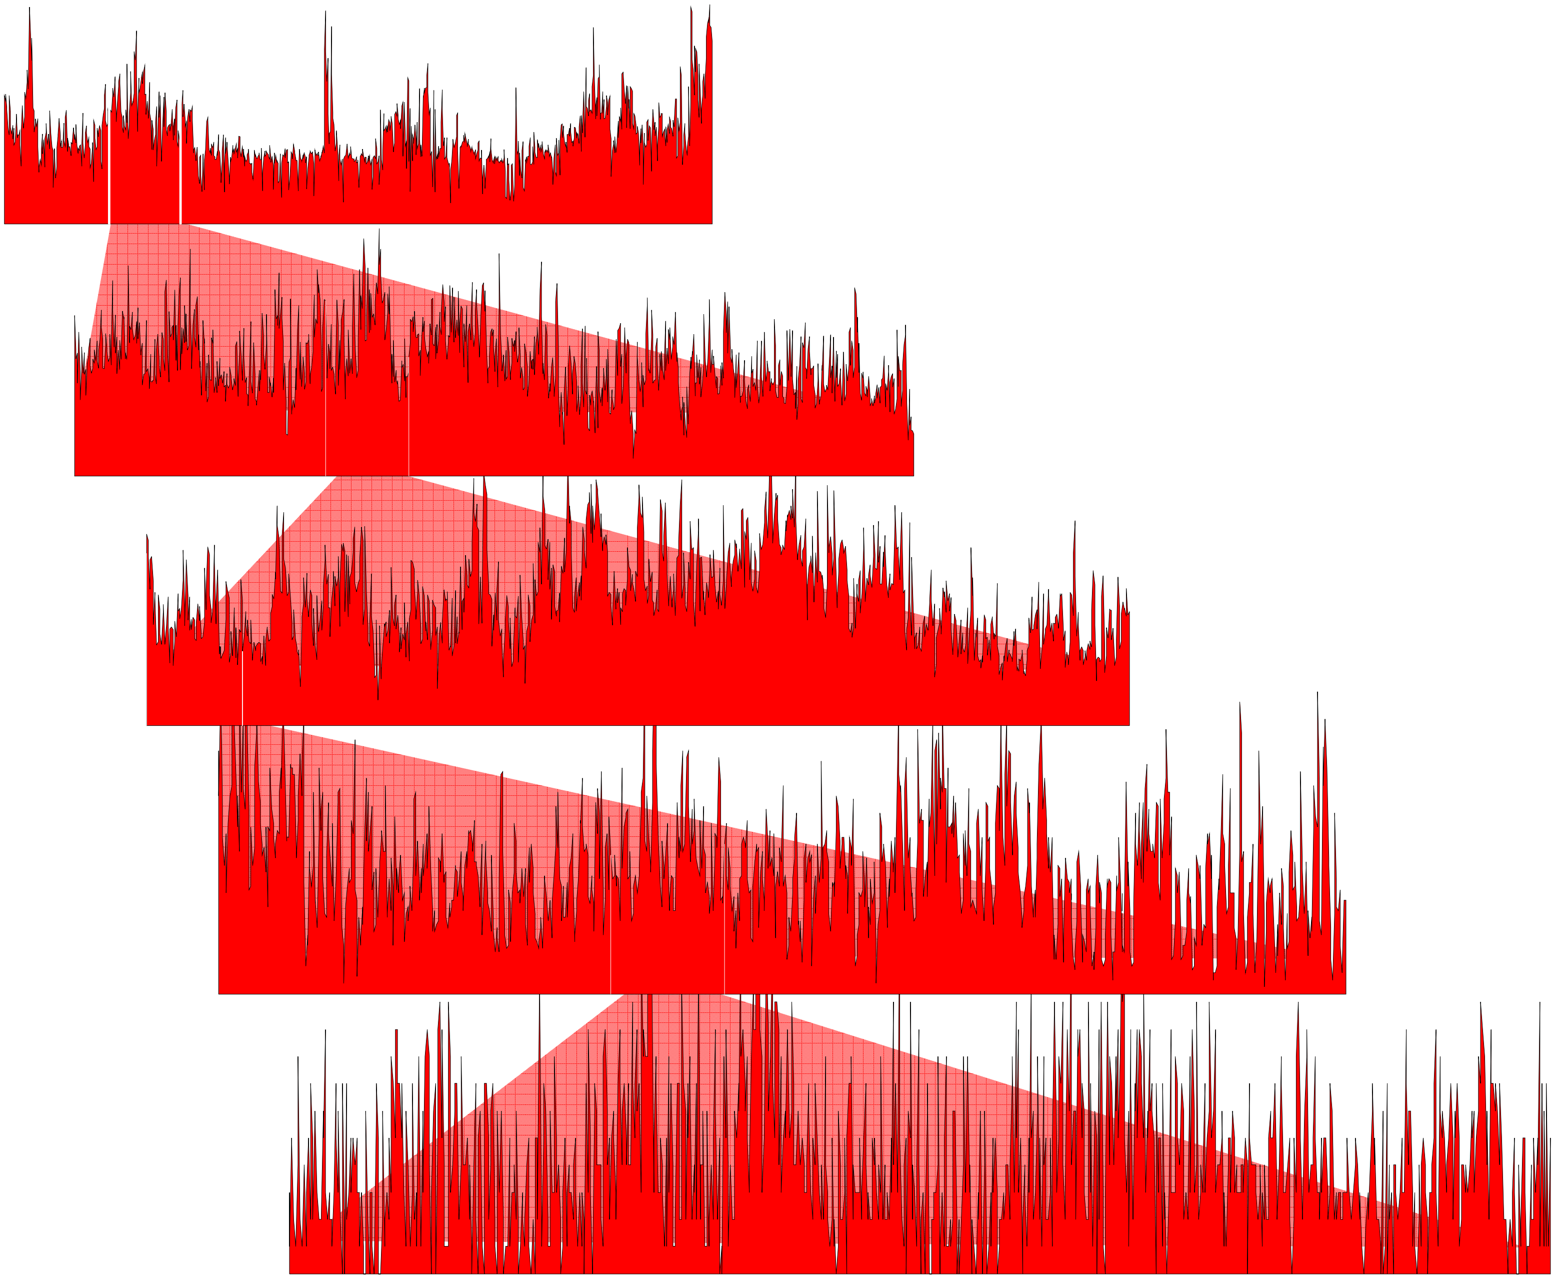

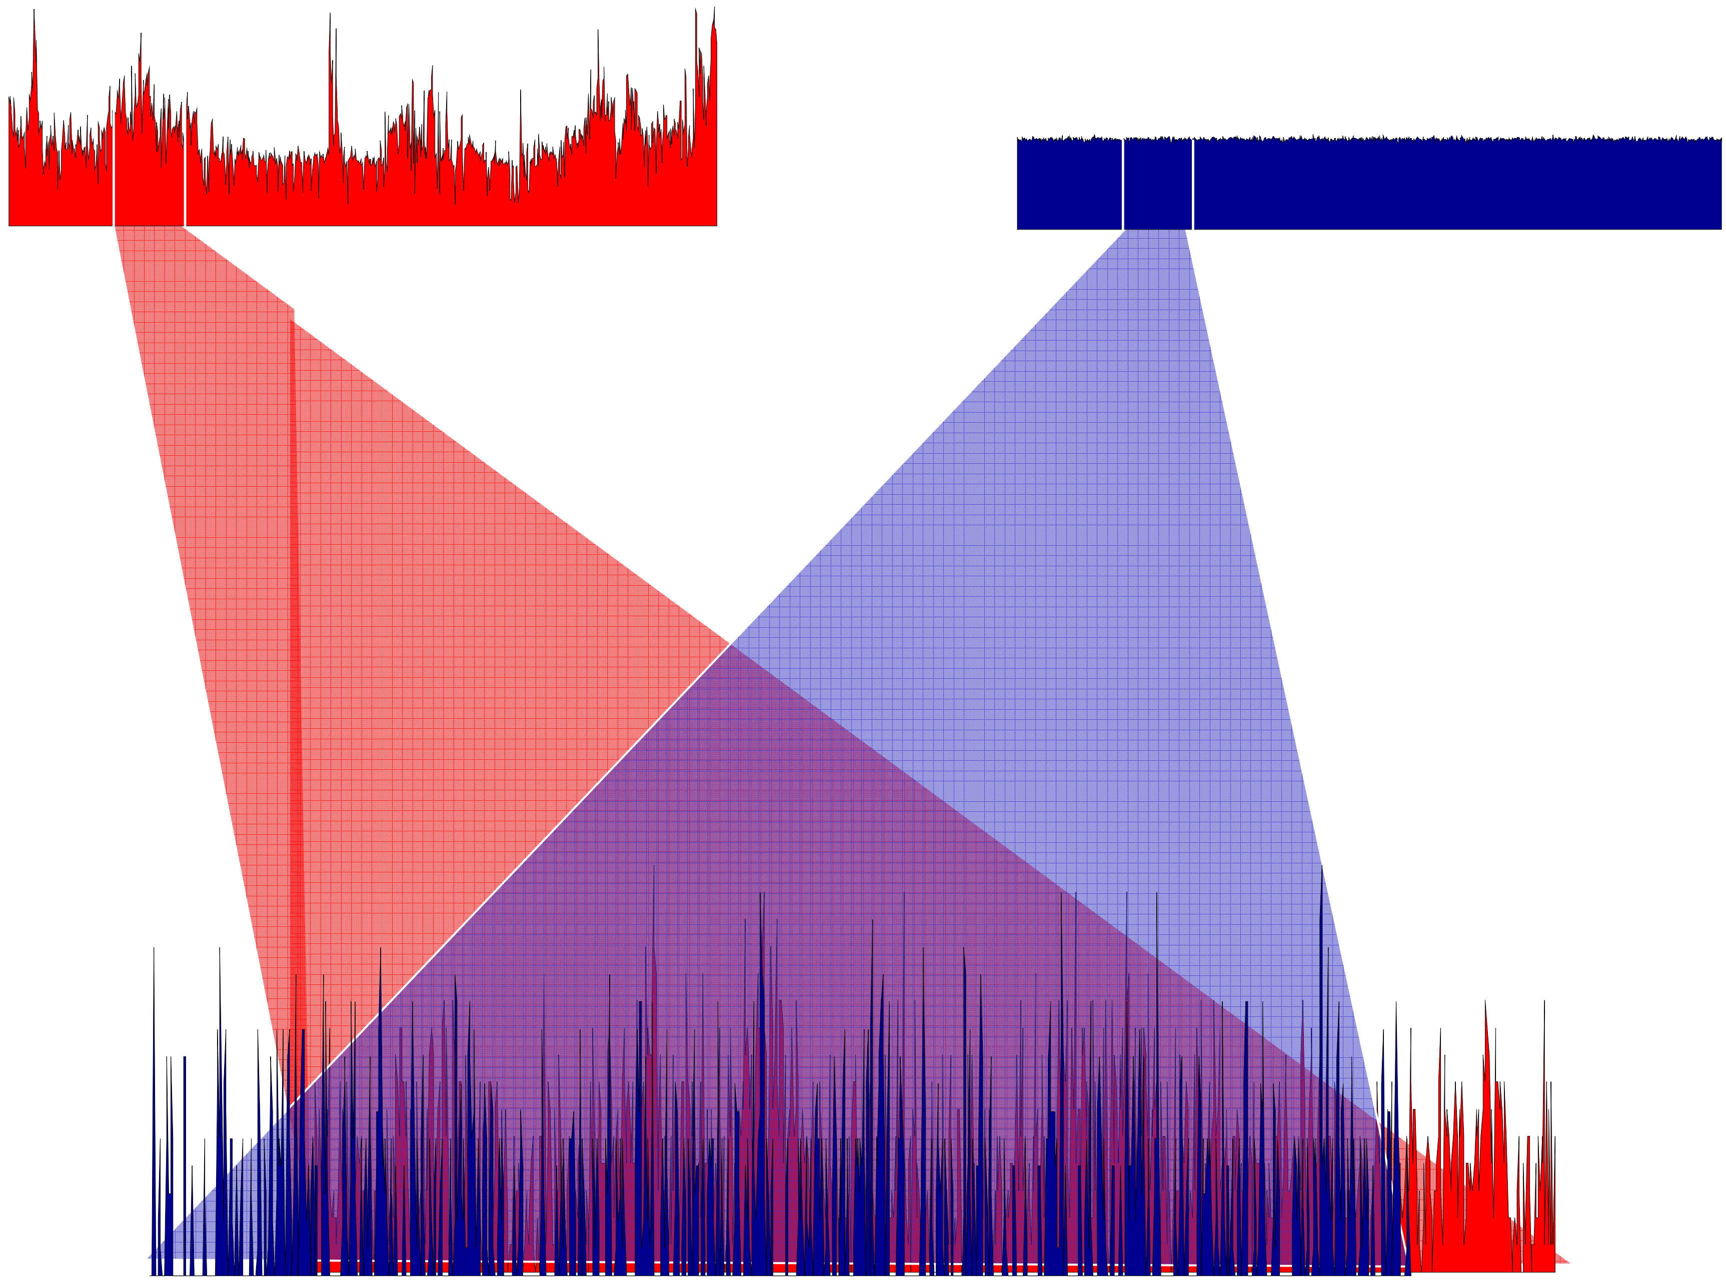

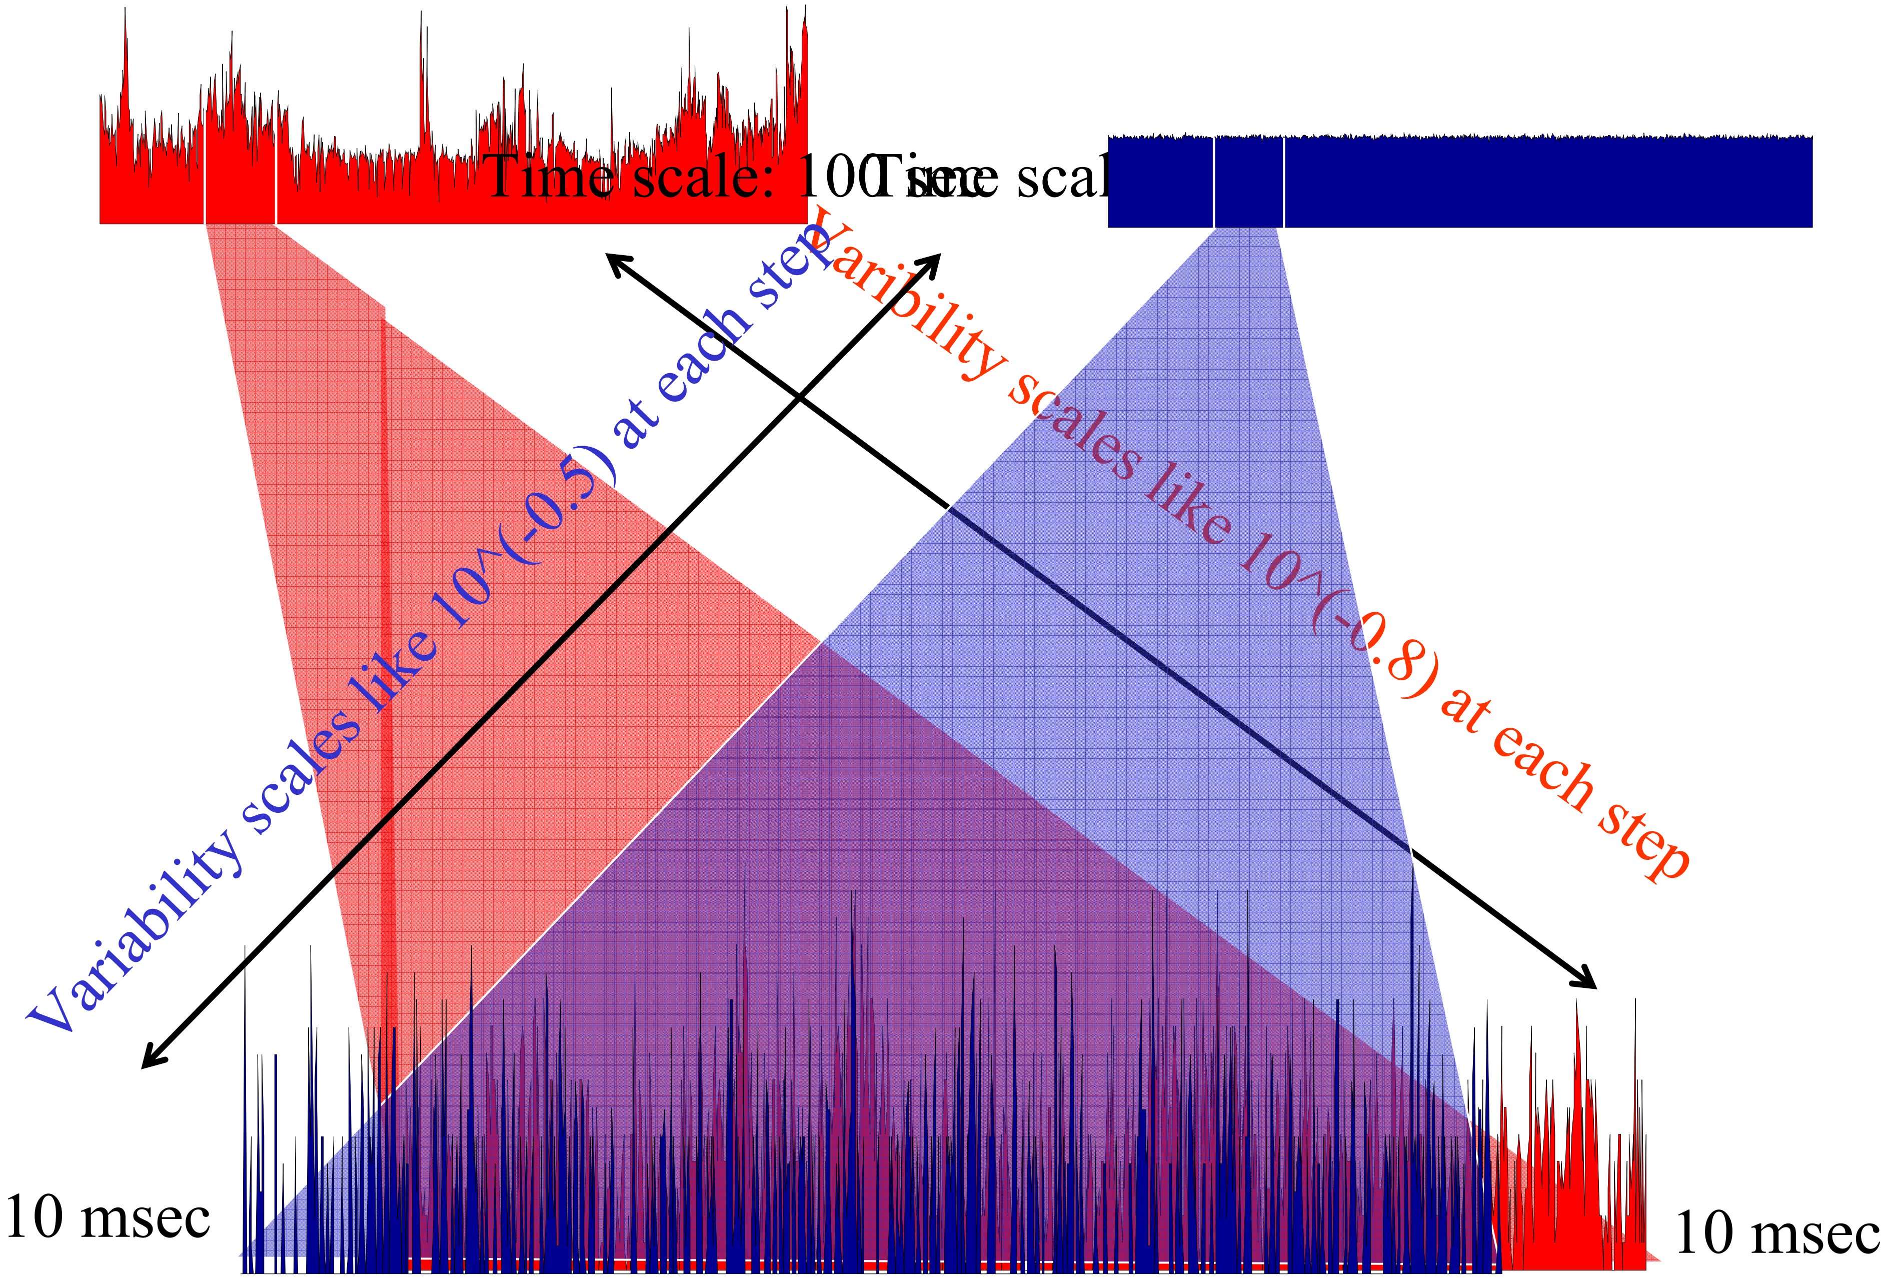

Time scale: 100 sec
Time scal
Variability scales like 10^(-0.5) at each step
Variability scales like 10^(-0.8) at each step
10 msec
10 msec

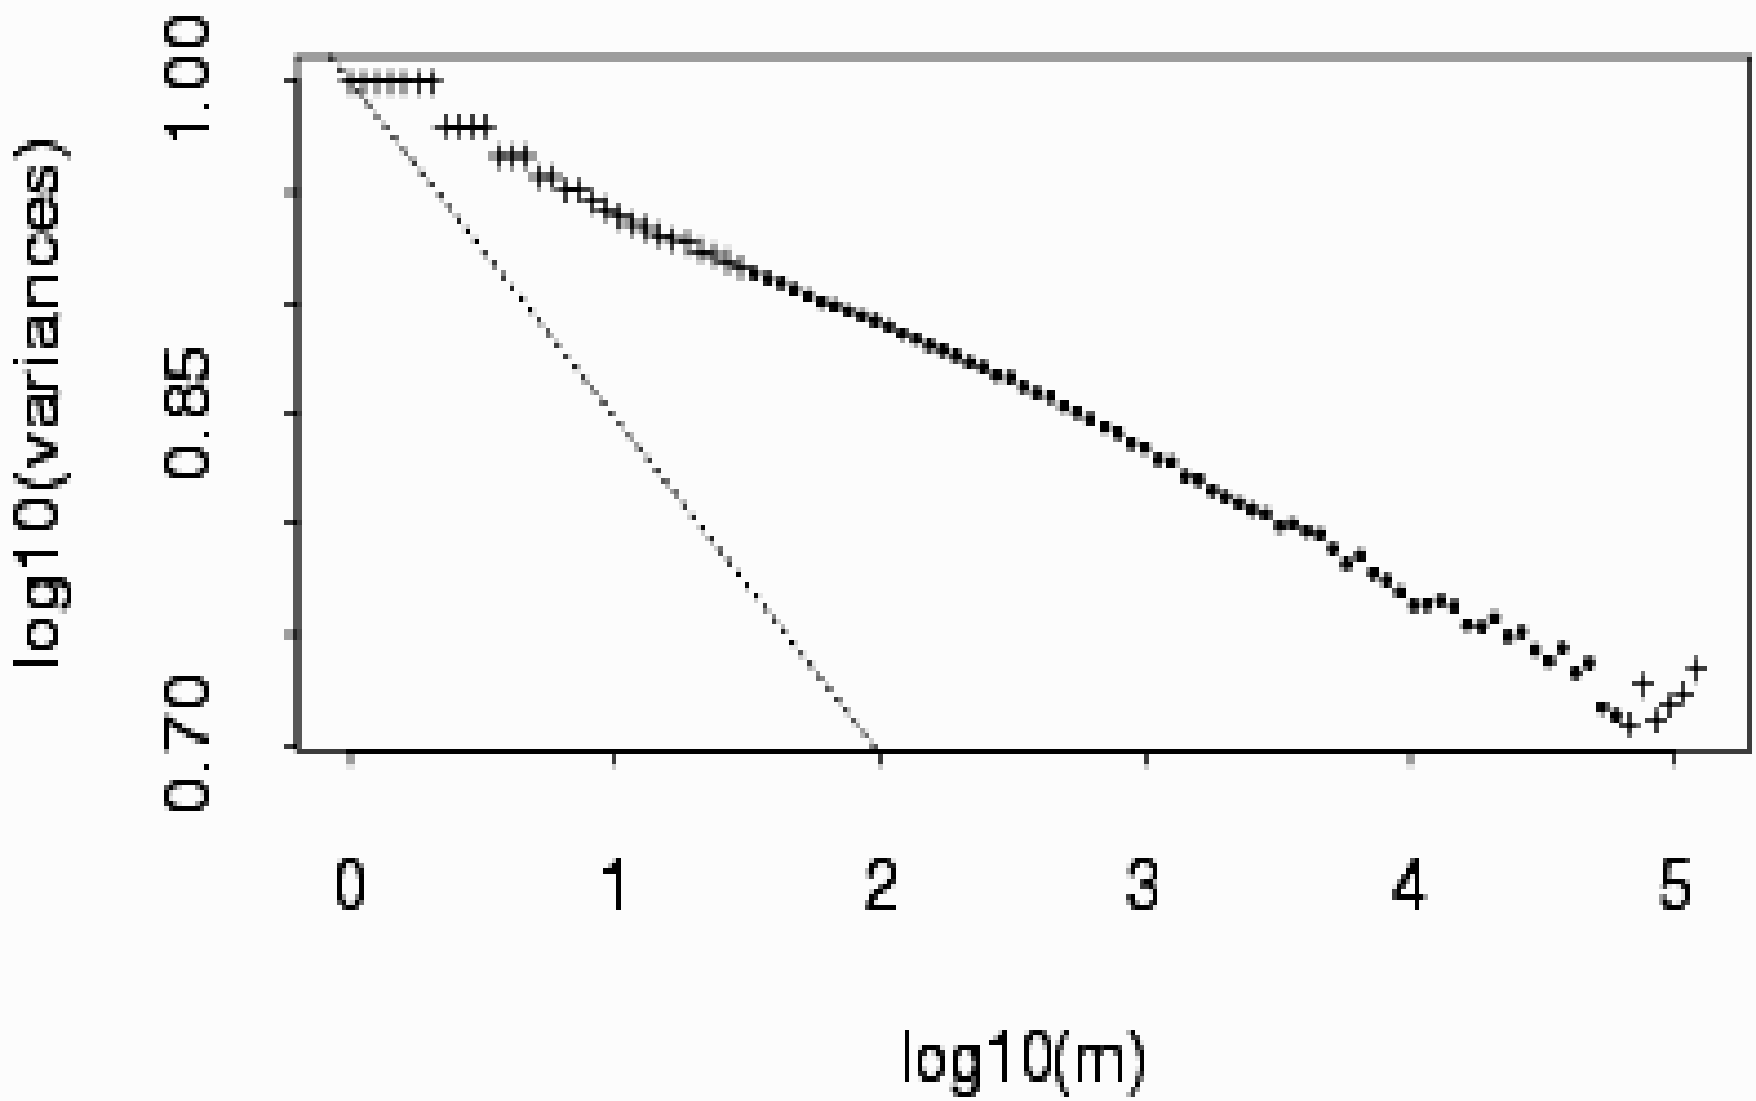
1.00
0.85
0.70
log10(variances)
0
1
2
3
4
5
log10(m)

# The "Physics" View: Self-Organized Criticality (SOC)

- Csabi, "*1/f noise in computer network traffic*" (1994)
- Ohira and Sawatari, "*Phase transition in computer traffic models*" (1998)
- Yuan, Ren, and Shan, "*Self-organized criticality in a computer network model*" (2000)
- Sole and Valverde, "*Information transfer and phase transitions in a model of Internet traffic*" (2001)

# SOC and Internet Traffic

- Can ignore Internet-specific “details” for the sake of addressing “fundamental” issues
- Self-similar Internet traffic is the result of a “critical” phase transition
- At criticality: 1/f-type traffic fluctuations, heavy-tailed queues and latencies, etc.
- There exist deep links between Internet traffic and highway traffic

# The SOC-type Network Model

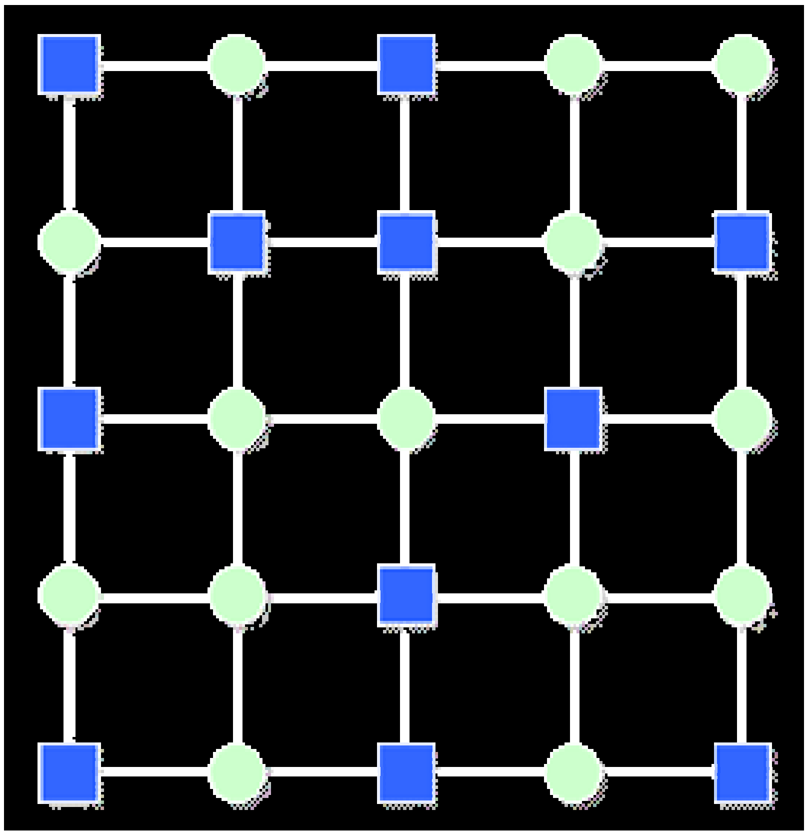

"SOCnet"

- 2-d lattice model with 4 nearest neighbors
- Some nodes are "hosts", some are "routers"
- Hosts generate packets at rate λ (pick random dest.)
- Routers store or forward packets (rule-based)
- Unbounded queues at the routers
- No congestion control!!

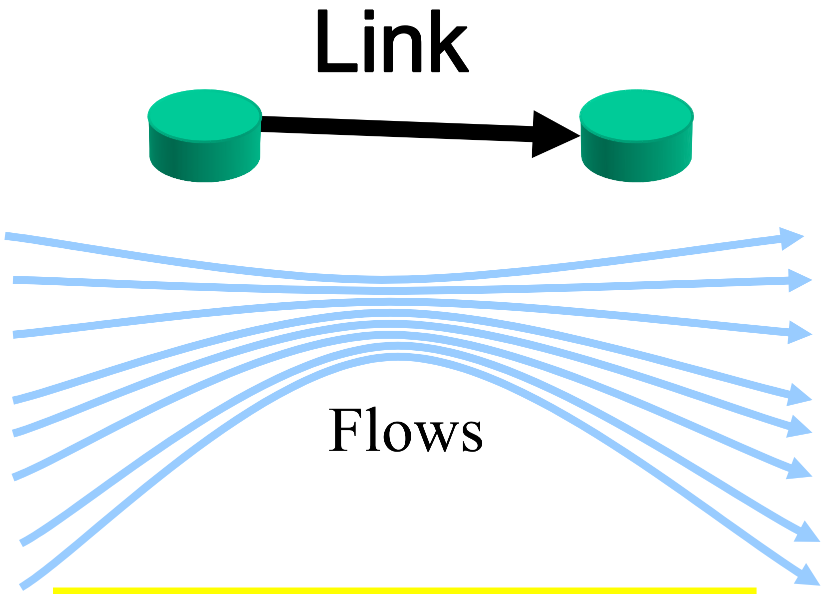


At each link

$$\sum_{k \in K_l} x_k^l \le c_l$$

flow ≤ capacity

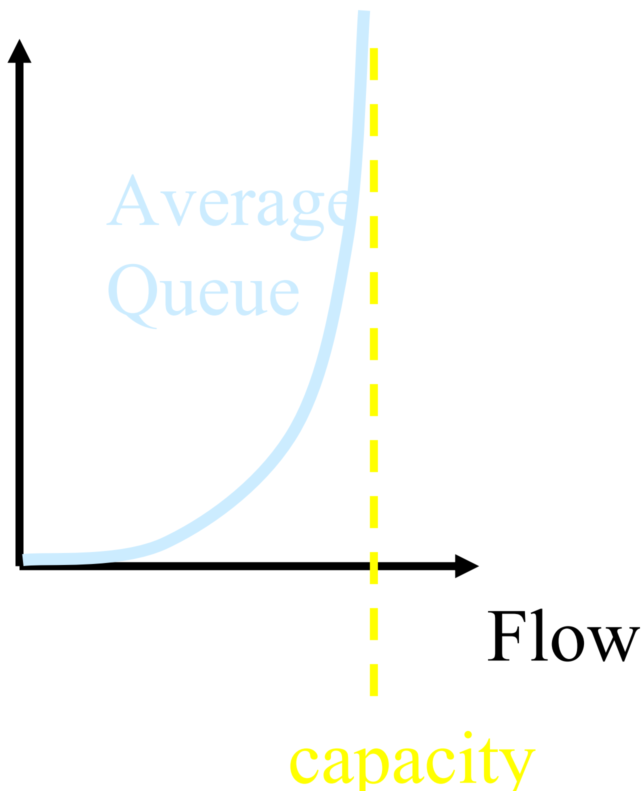

# SOCnet and Criticality

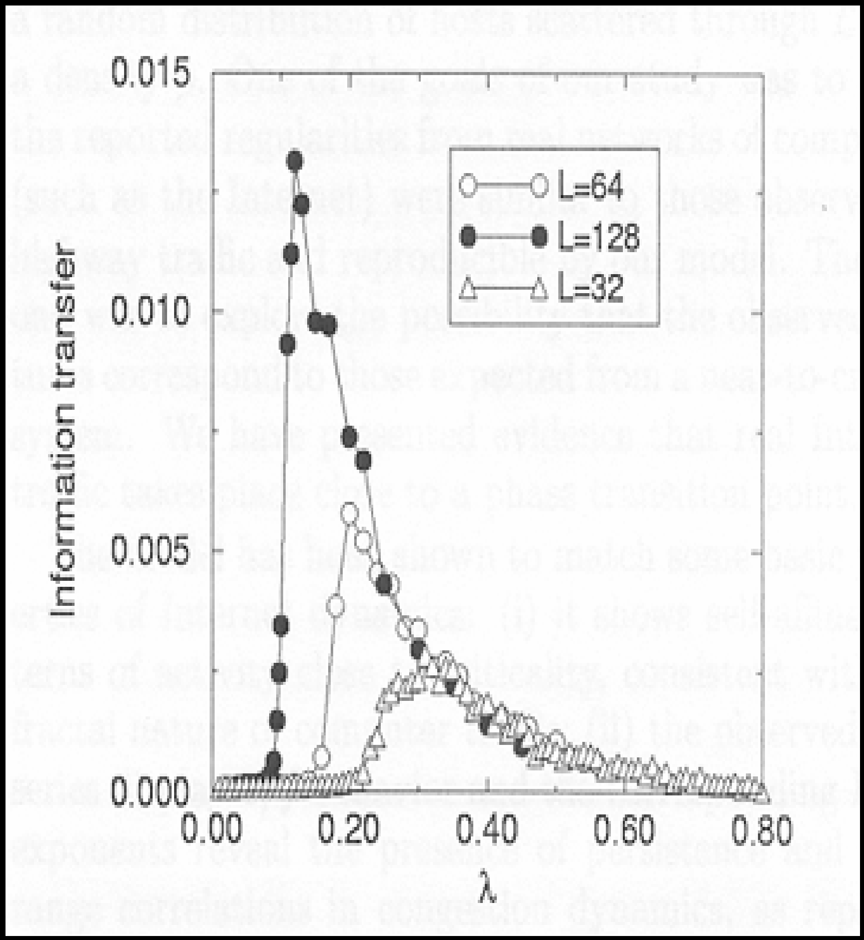


Sole and Valverde(2001)

- Phase transition at $\lambda = \lambda^*$
  - Average latency bifurcates
  - Packet flow is maximized
- At criticality ($\lambda = \lambda^*$)
  - Self-similar fluctuations
  - Power law latencies and queue lengths
  - Maximum efficiency, information transfer

# SOCnet and the Internet

- There exists no “critical” rate!
- Self-similar traffic is observed for all load levels
- Bottom line
  - It is conceivable to build a SOCnet
  - No engineer would ever design a SOCnet
  - No user would ever get anything done using SOCnet
  - Chaos-based explanations for self-similar traffic have been proposed and are equally specious

# The Networking View: High Variability I

- Leland, Taqqu, Willinger, and Wilson, "On the self-similar nature of Ethernet traffic" (1993)
- Paxson and Floyd, "Wide area traffic: The failure of Poisson modeling" (1994)
- Crovella and Bestavros, "Self-similarity in World-Wide-Web traffic" (1996)
- Mandelbrot, "Long-run linearity, locally Gaussian processes, h-spectra, and infinite variances (1969)

*Application* traffic (e.g., WWW)

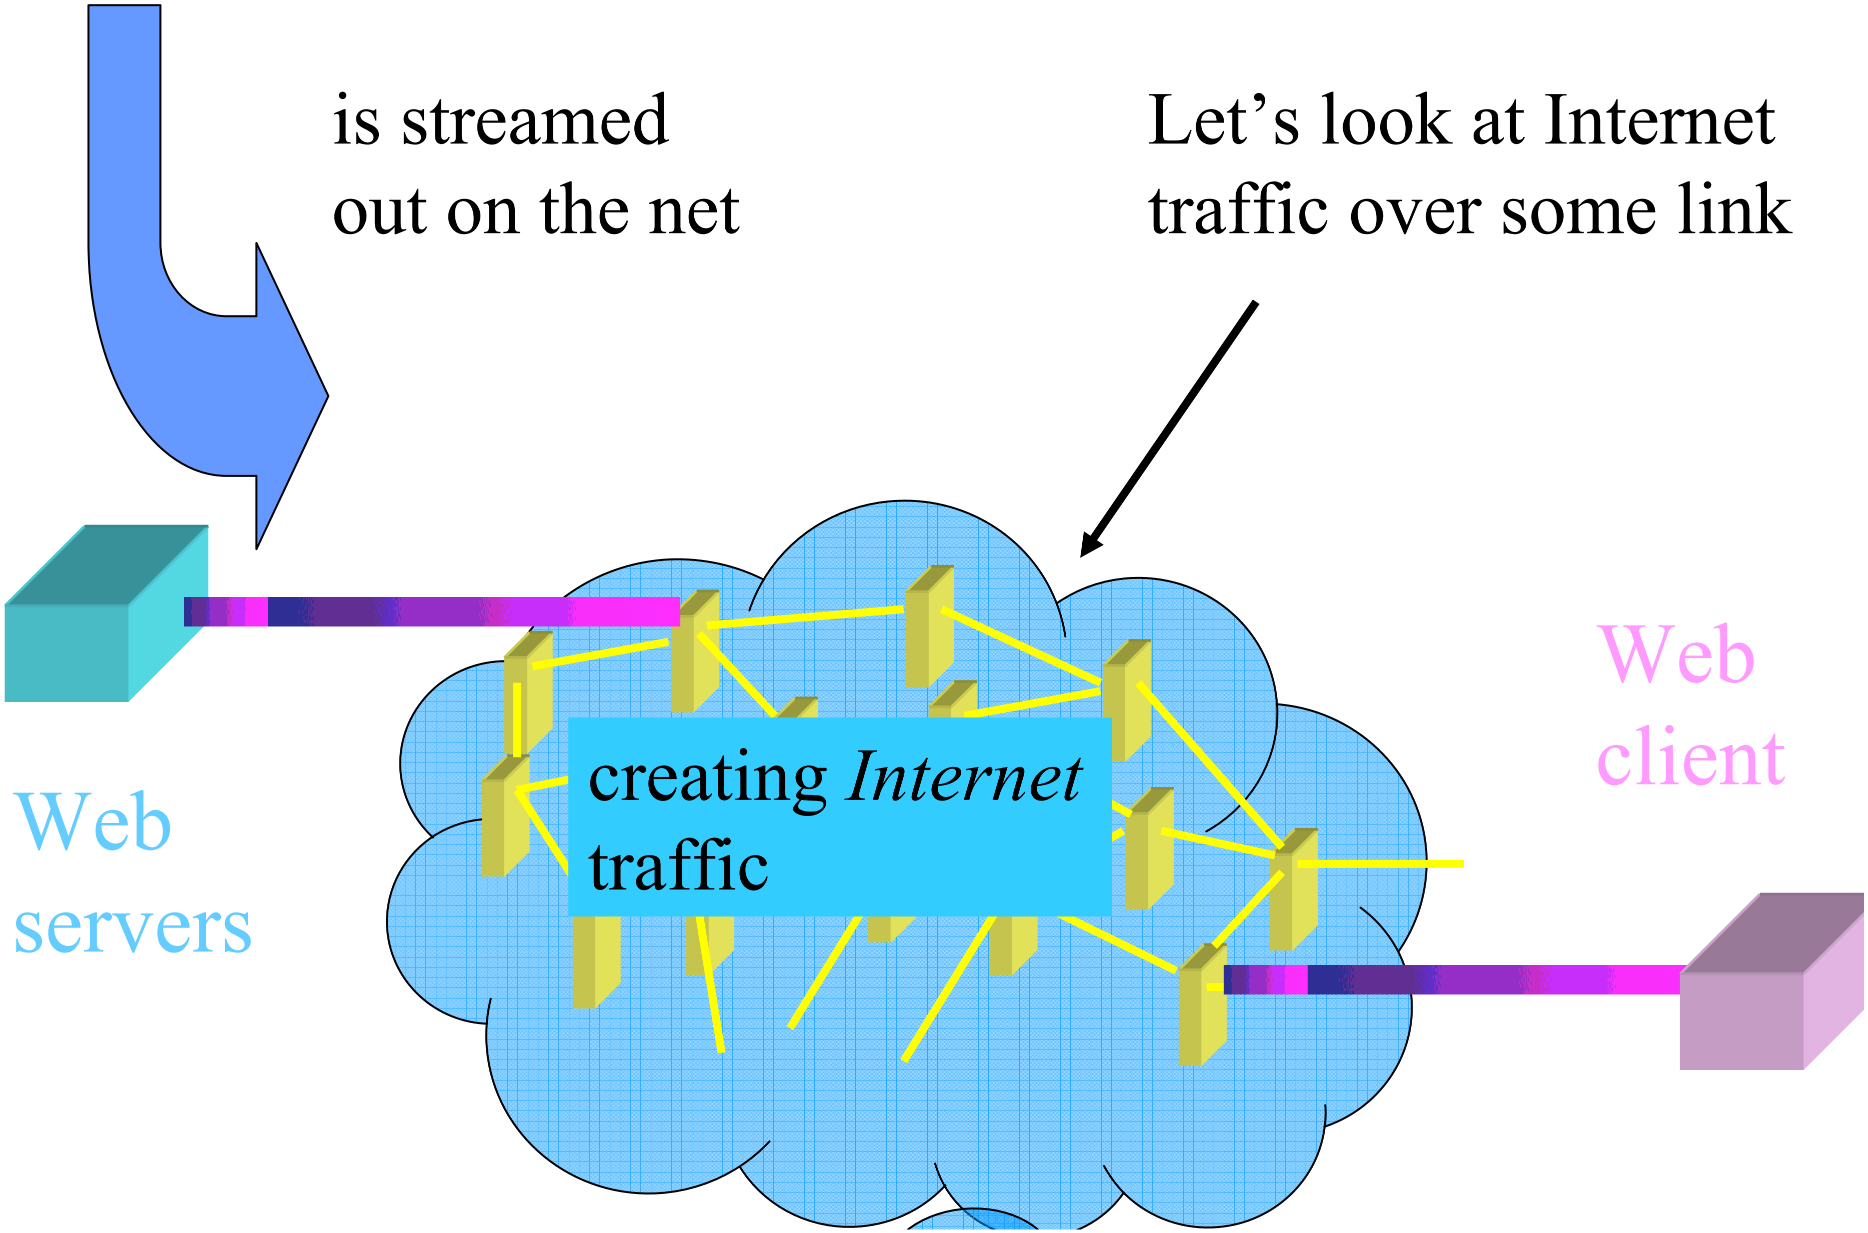

# Internet traffic

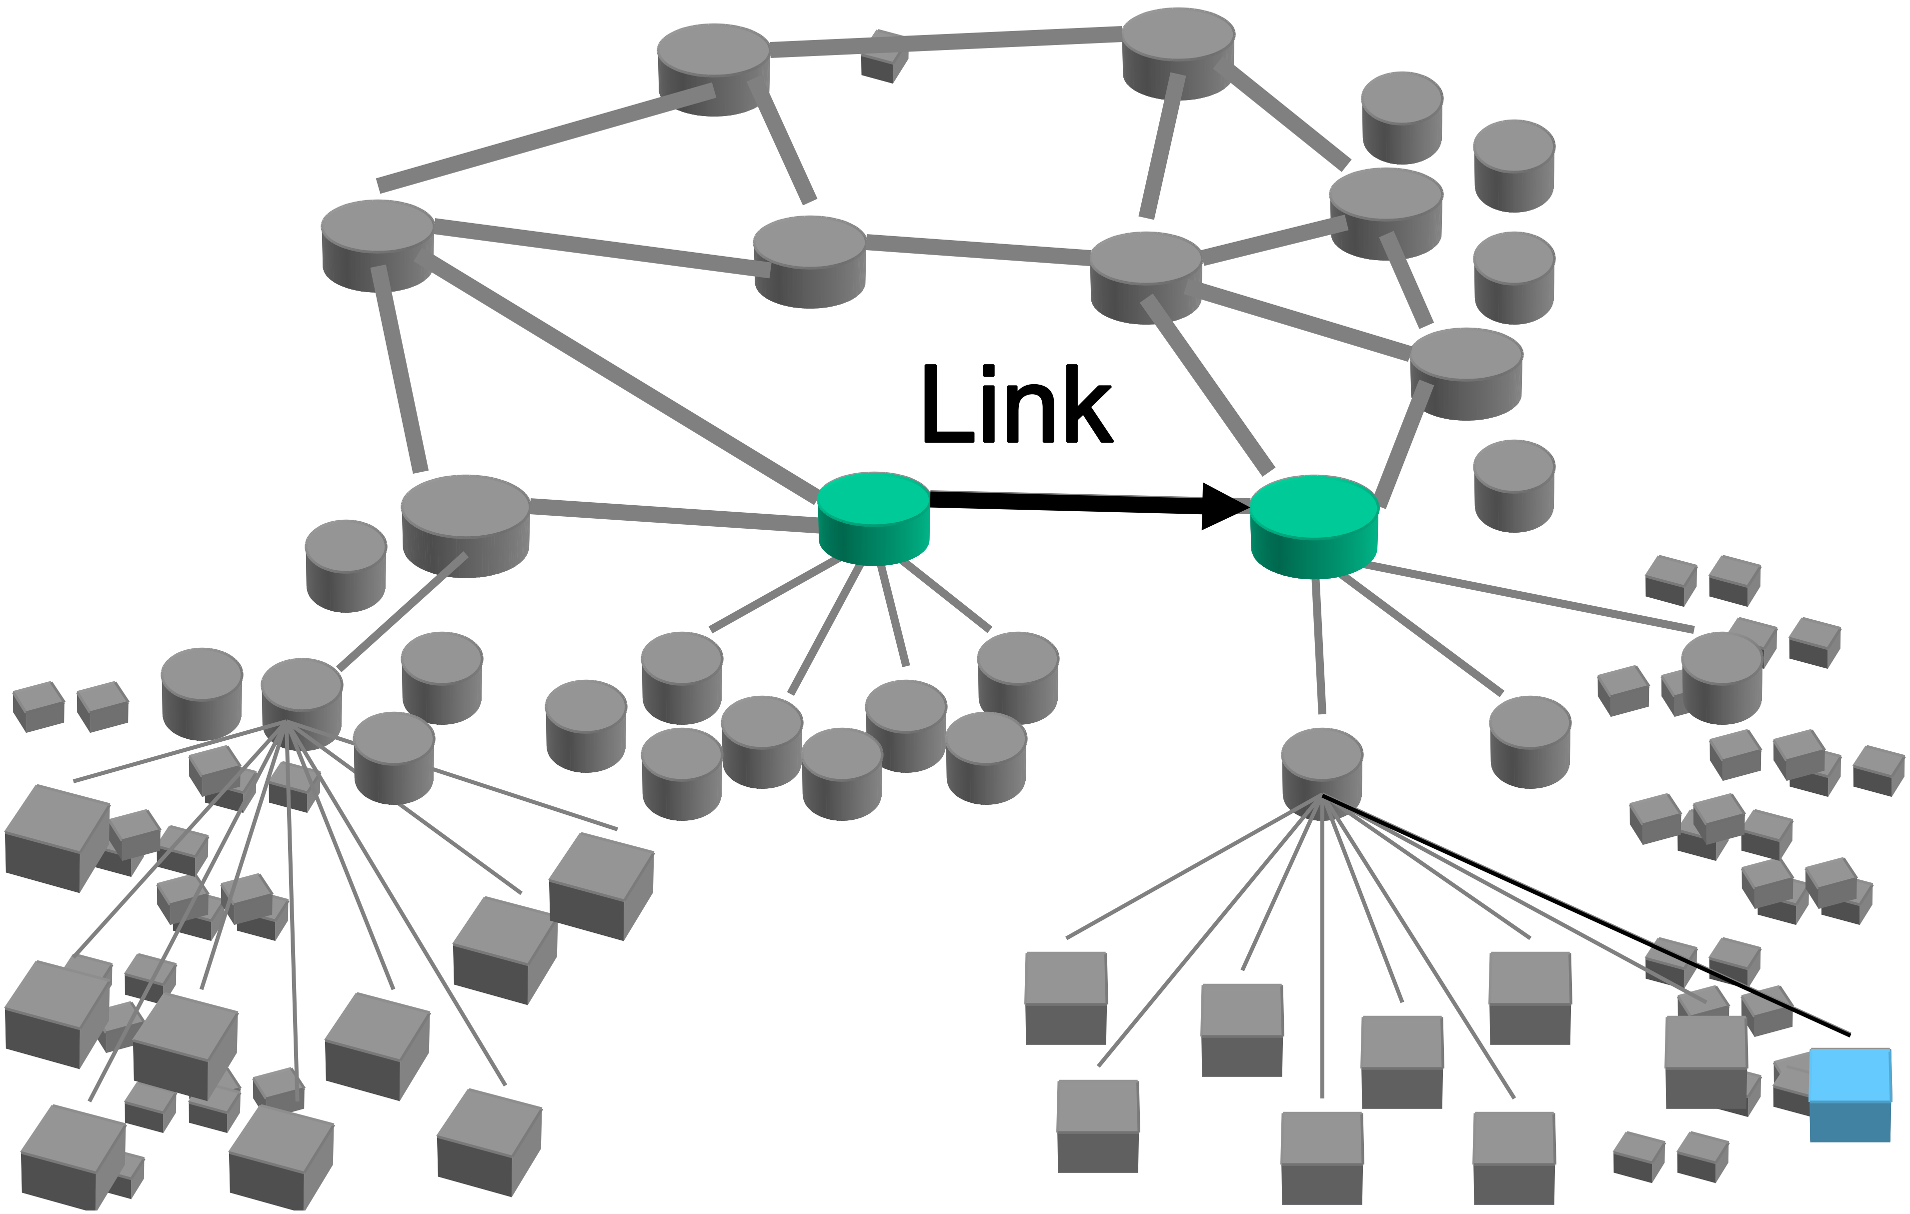

Web traffic …

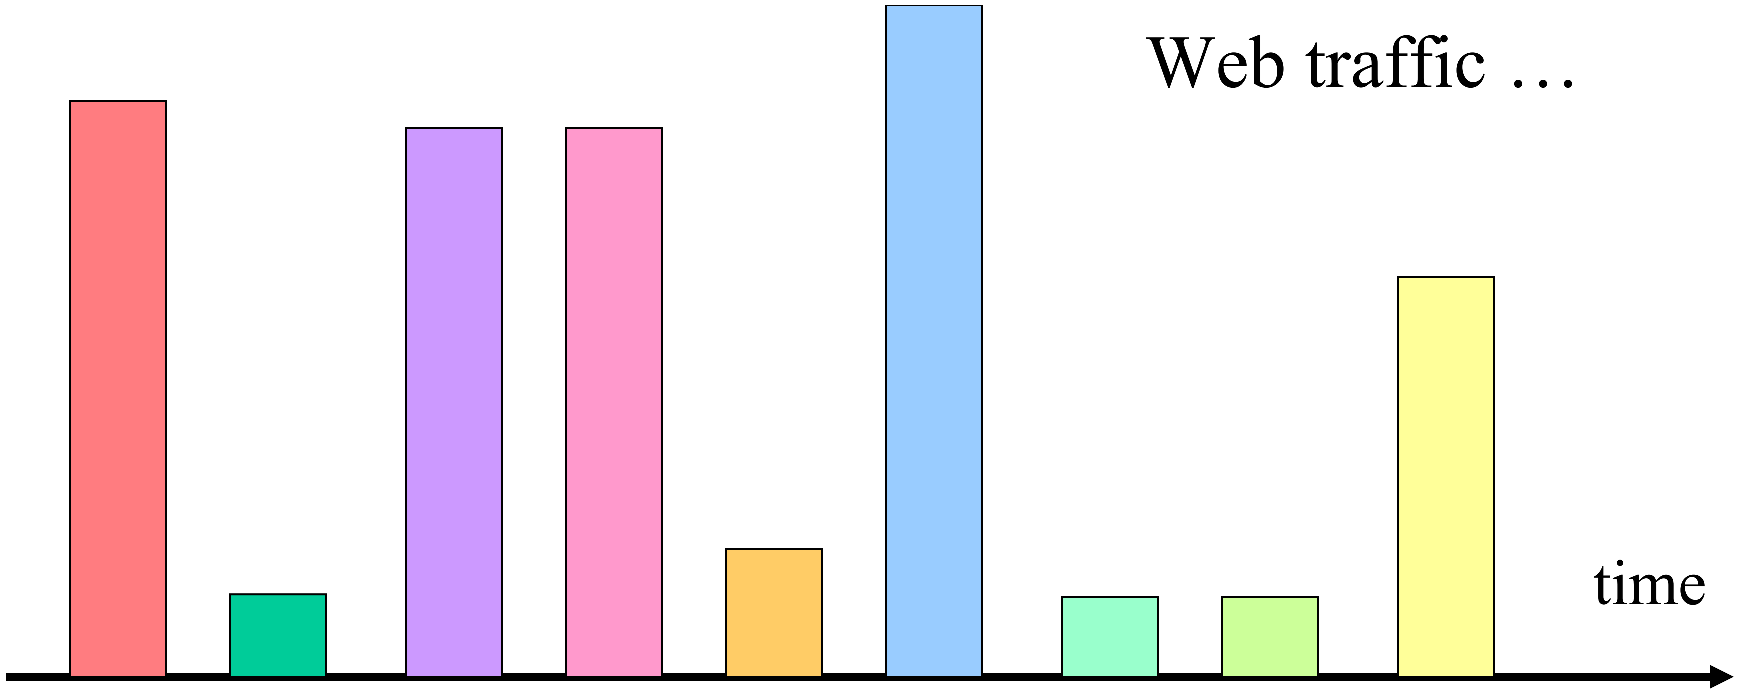


… is streamed onto the Internet …

… creating “bursty-looking” link traffic

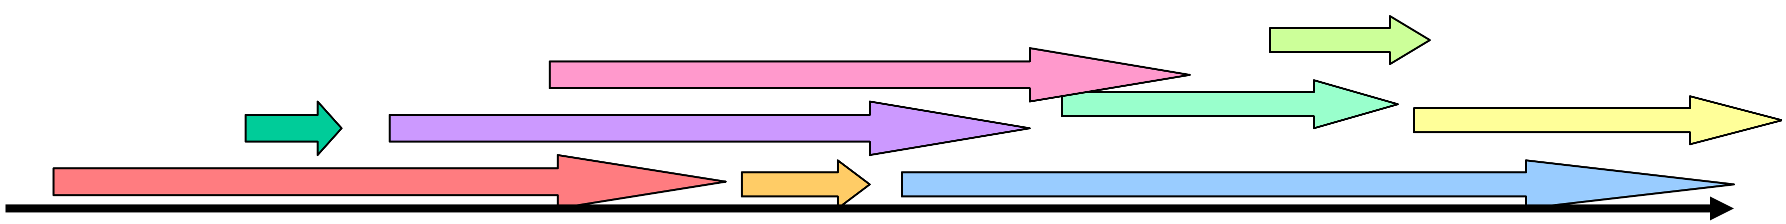

# Mandelbrot-Cox Construction

- Renewal reward processes and their aggregates
  - Aggregate is made up of many constituents/flows
  - Flows arrive at random
  - Each flow has a “size” (number of packets)
  - Each flow transmits it packets at some (constant) rate
- Mathematical results
  - If flows arrive at random and if their sizes are heavy-tailed with infinite variance, then the aggregate traffic is perforce self-similar (exhibits long-range dependence)
  - Aggregate traffic is Fractional Gaussian noise

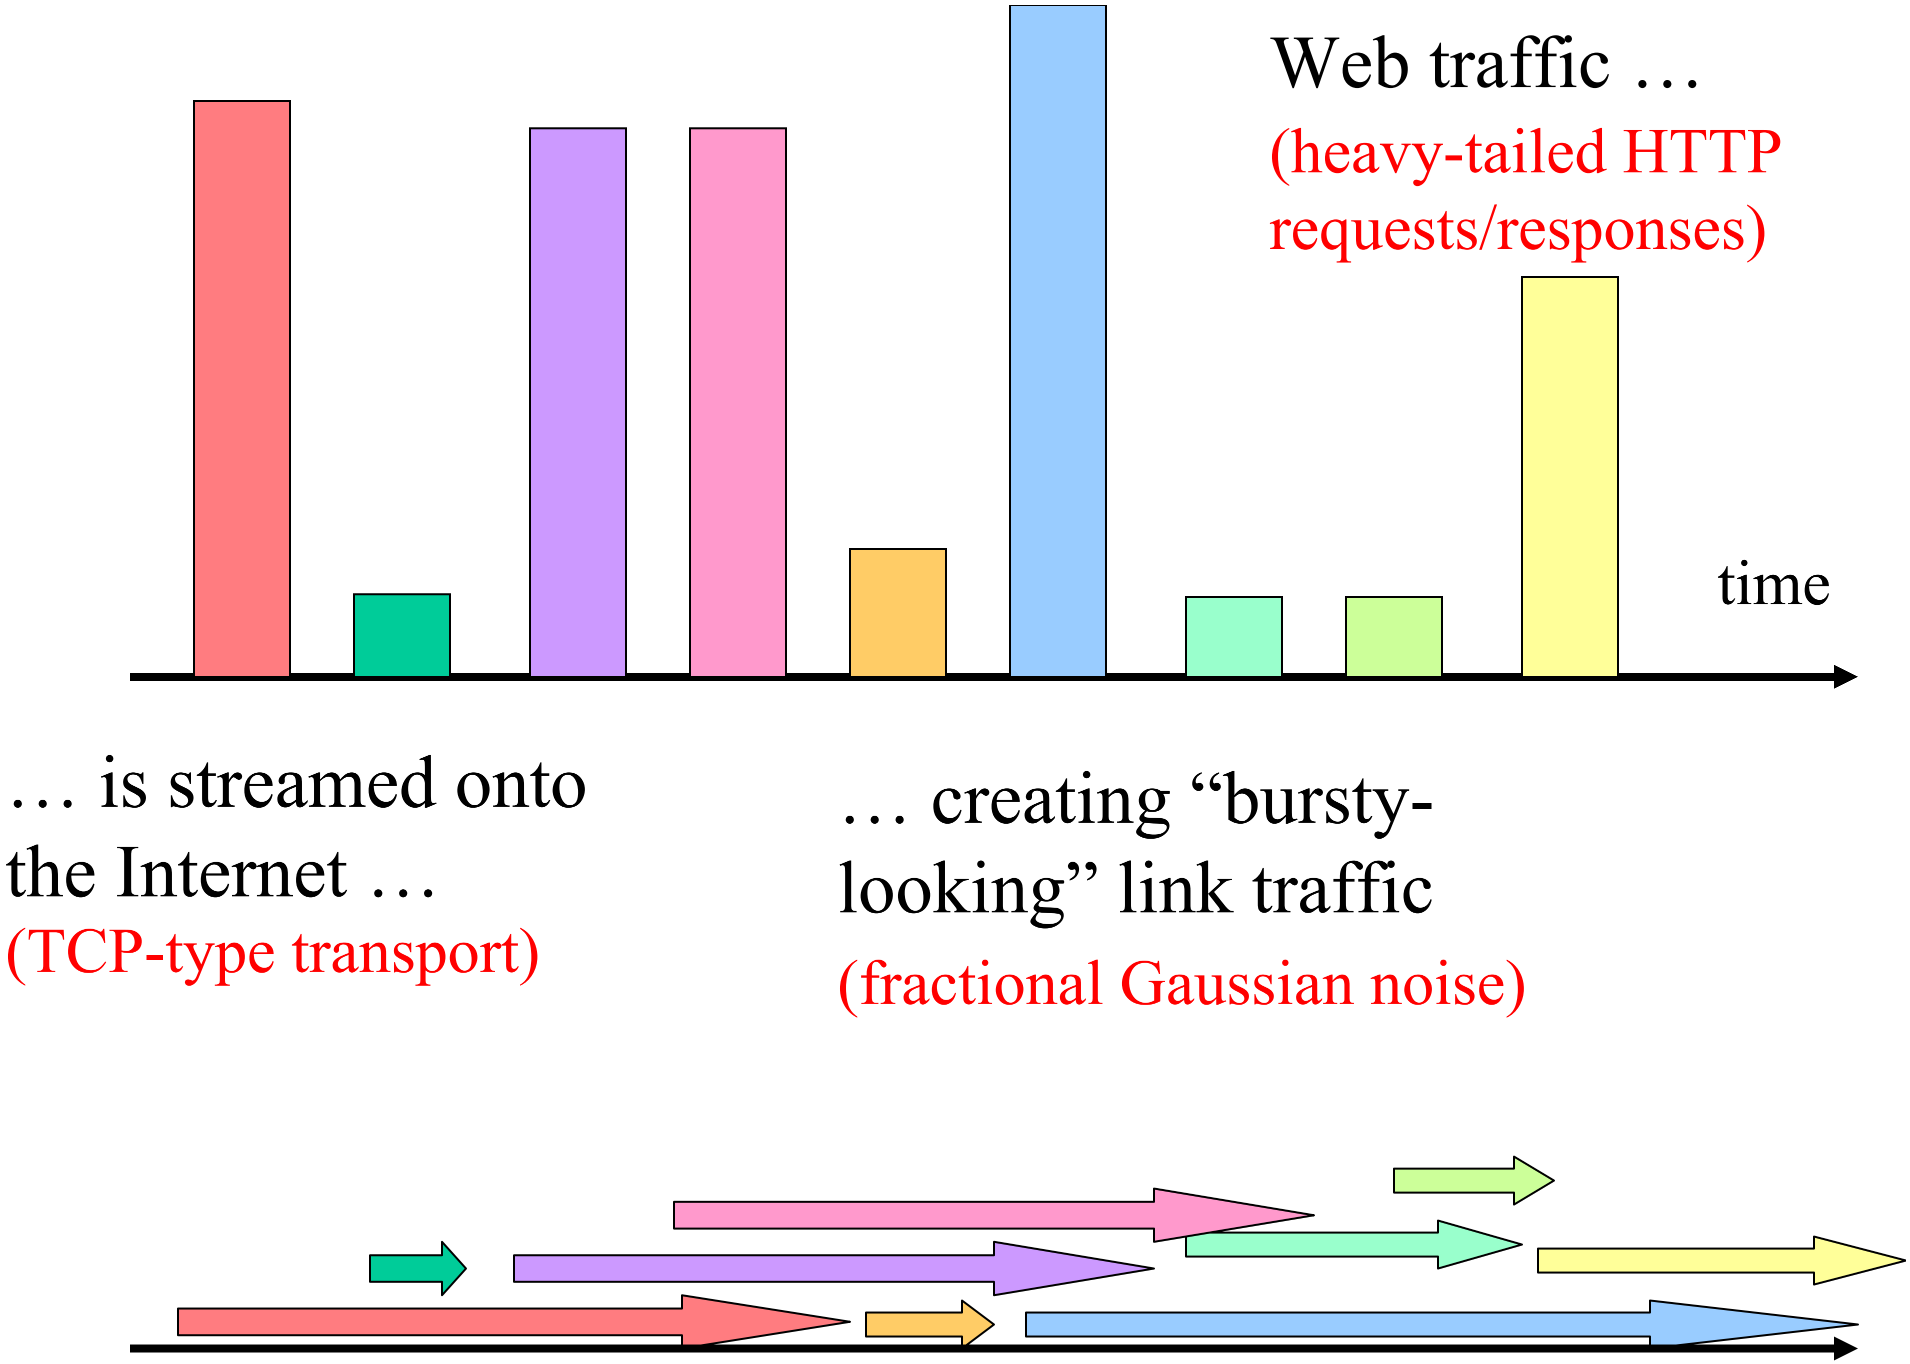
Web traffic …
(heavy-tailed HTTP requests/responses)
time
… is streamed onto the Internet …
(TCP-type transport)
… creating “bursty-looking” link traffic
(fractional Gaussian noise)

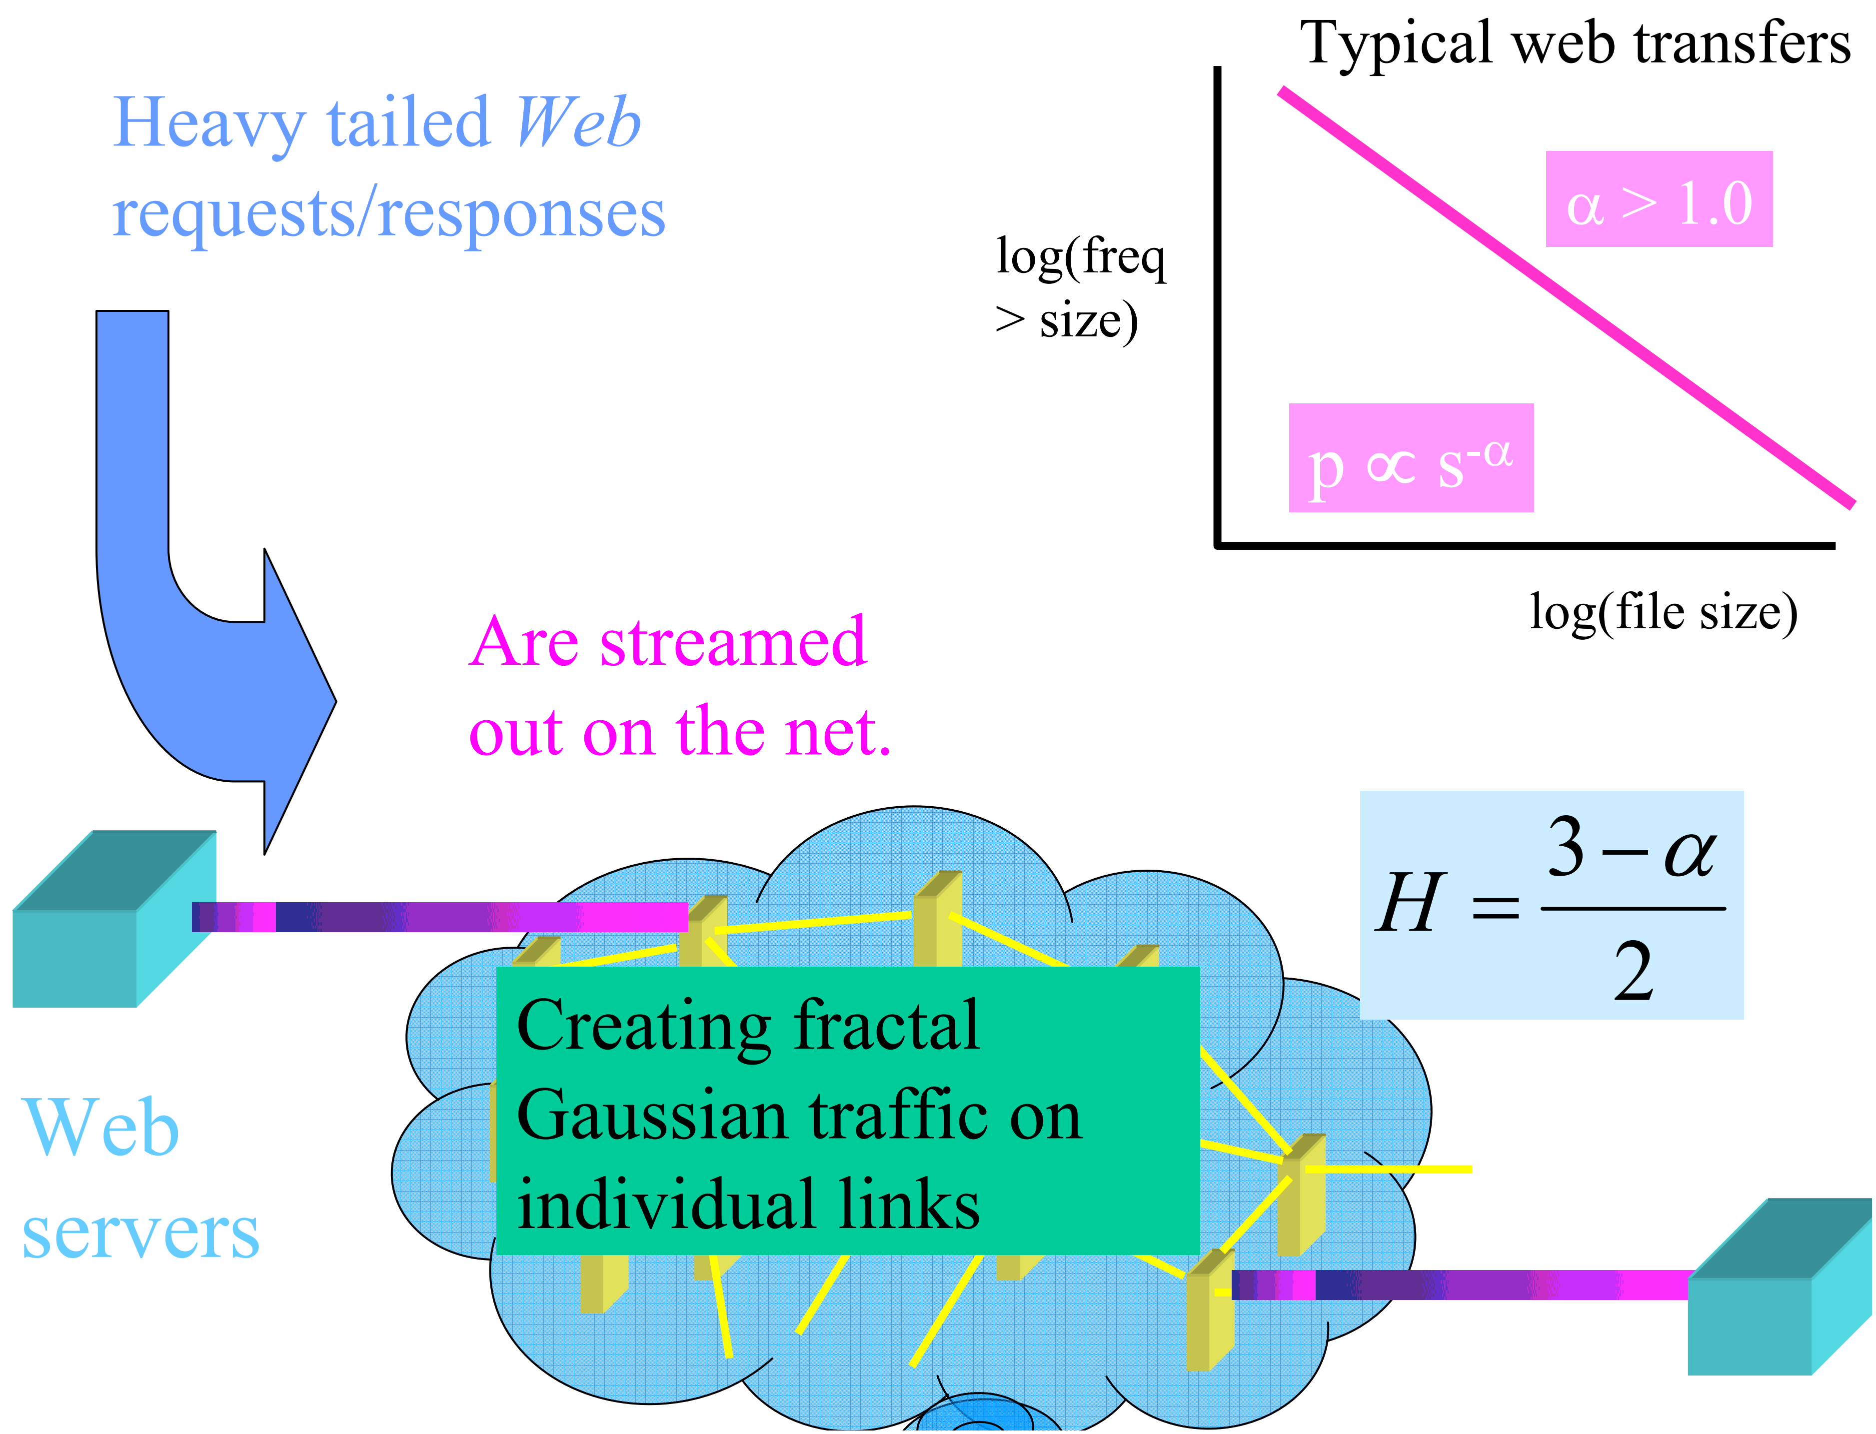
Heavy tailed *Web* requests/responses
Typical web transfers
log(freq > size)
log(file size)
α > 1.0
p ∝ s^-α
Are streamed out on the net.
Web servers
Creating fractal Gaussian traffic on individual links
H = (3 − α) / 2

# Measured IP Flow Property (Size)

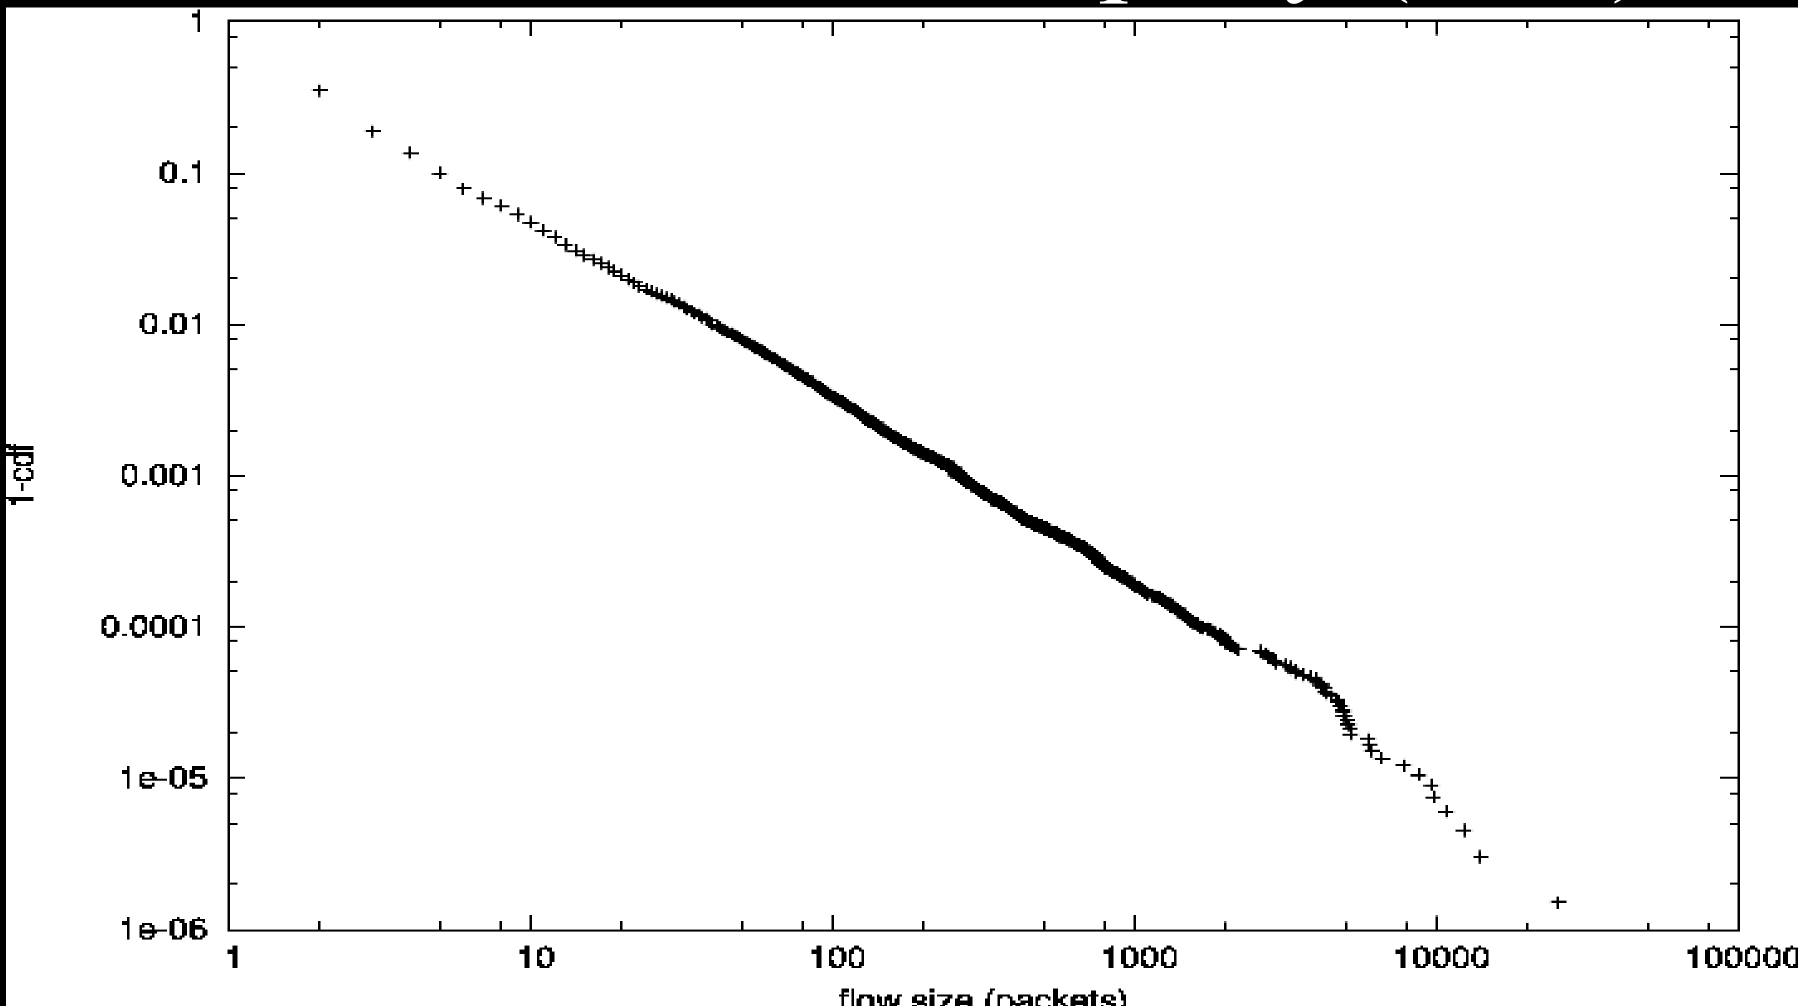

# High Variability and Internet Traffic

- Coherent and mathematically rigorous framework for Internet traffic (large time scales)
- Flow measurements over the last 10 years have consistently shown heavy-tailed/infinite variance behavior
- Prime example for what is meant by a consistent, mathematically rigorous, networking-based, explanation with supporting measurements
- Says nothing about Internet traffic over short time scales

# 2001: JDC's 1st INFOCOM Paper

Zhu, Yu, and Doyle, "Heavy tails, generalized coding, and optimal web layout", Proc. IEEE Infocom'01

- Provides an explanation of the origin of "heavy tails" in Web documents
- Suggests that heavy tails are here to stay
- Completes the explanation of self-similarity
- Motivates subsequent work on congestion control for "mice/elephant" traffic

# Topics Covered

- The Self-Similar Nature of Internet Traffic
  - The "Physics of Complex Systems" view
  - The Networking view
  - JCD's dream comes true: 1st INFOCOM paper
- High Variability in Internet Topology
  - The "Physics of Complex Systems" view
  - The Networking view
  - JCD's dream comes true: 1st SIGCOMM paper
- Lessons Learned

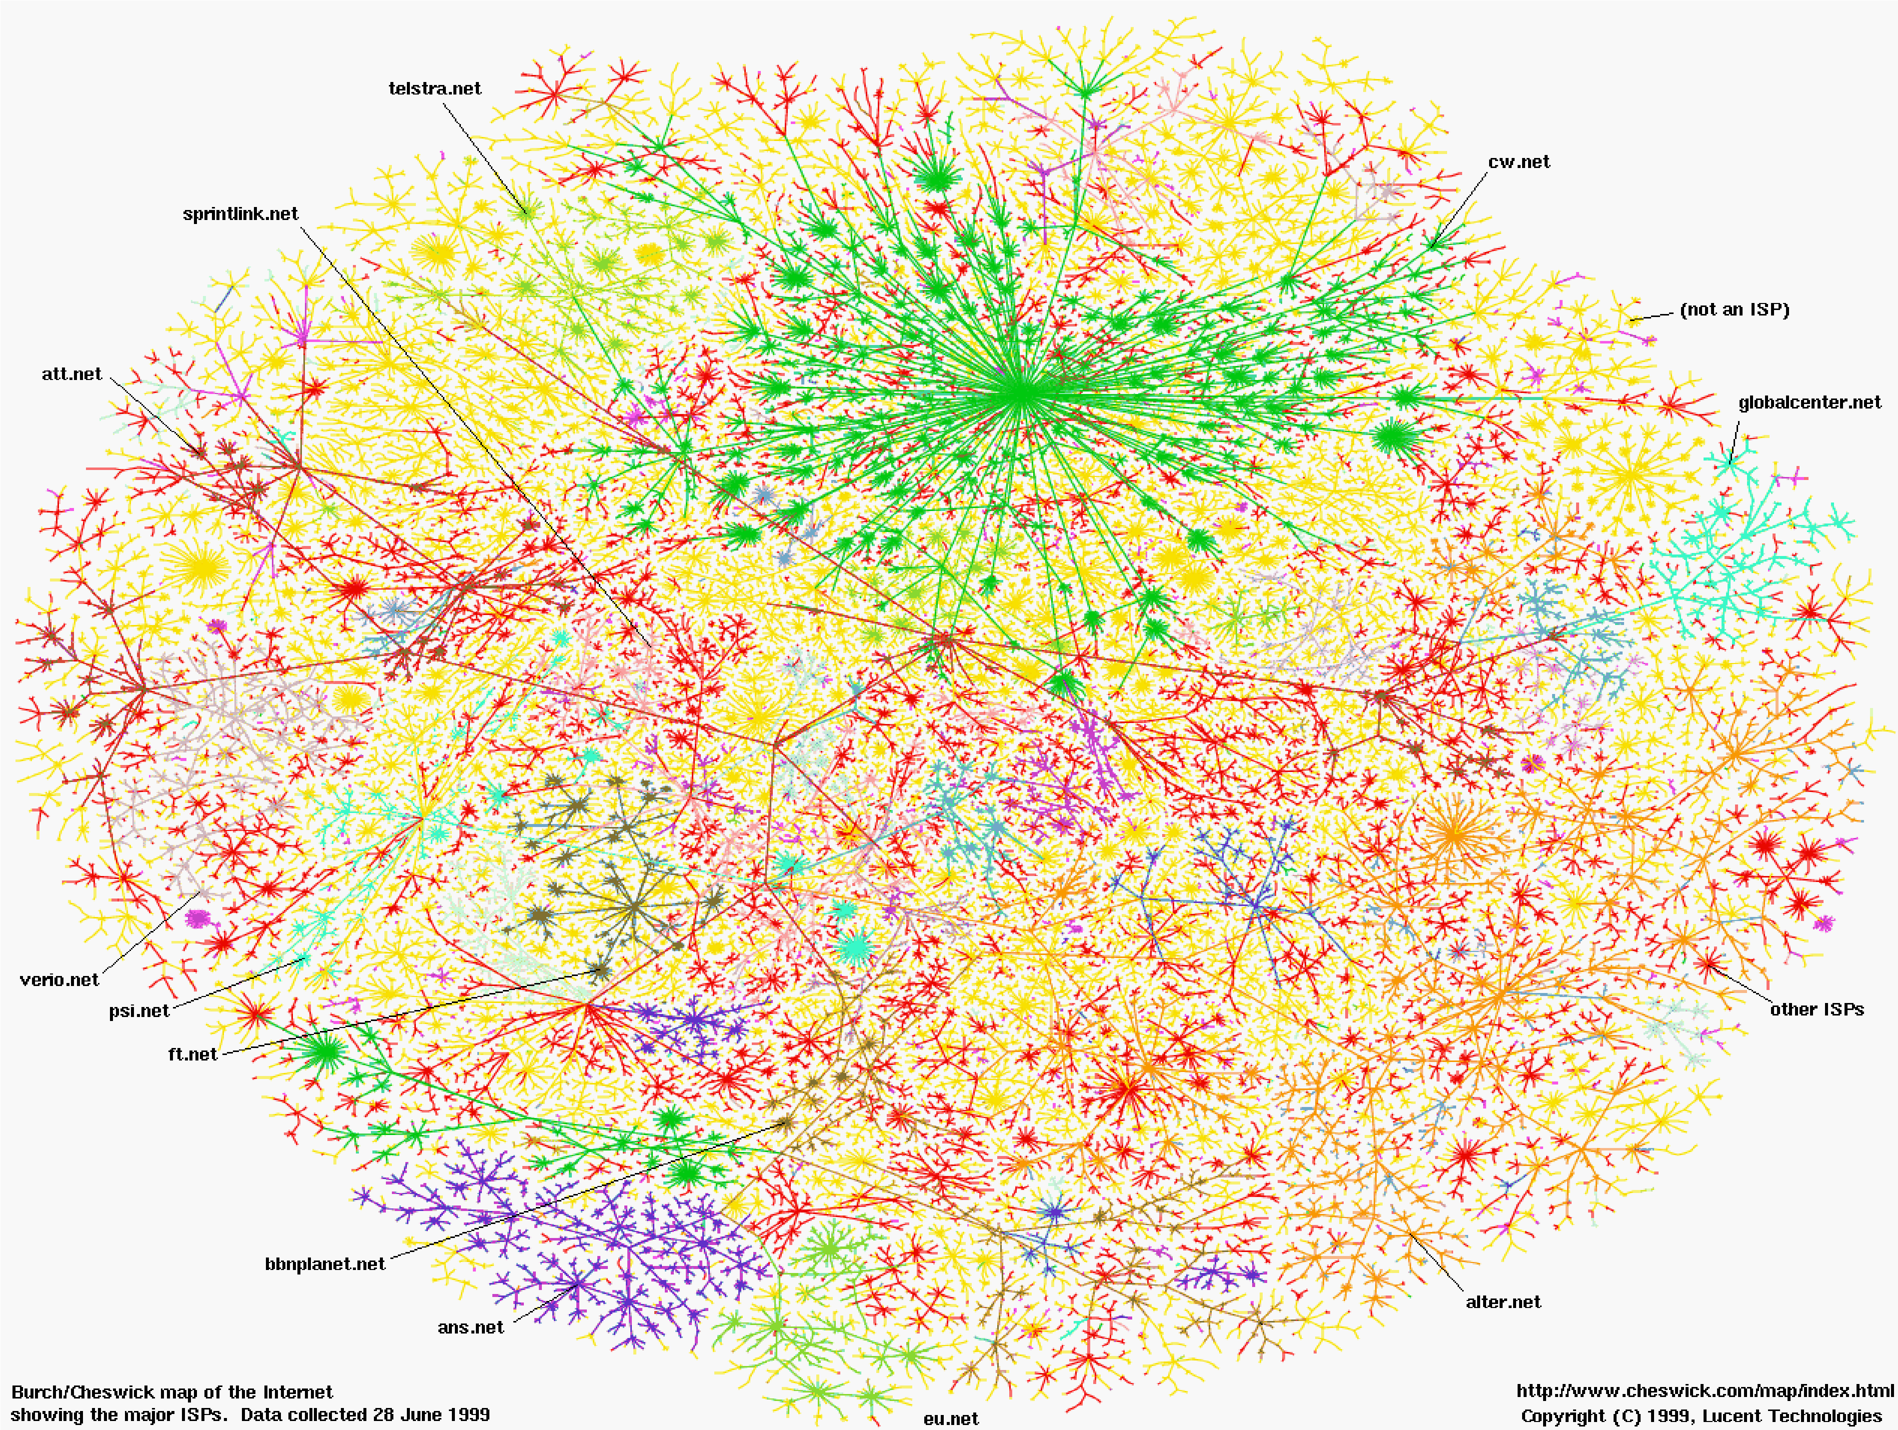
telstra.net
cw.net
sprintlink.net
(not an ISP)
att.net
globalcenter.net
verio.net
psi.net
ft.net
other ISPs
bbnplanet.net
ans.net
alter.net
eu.net
Burch/Cheswick map of the Internet
showing the major ISPs. Data collected 28 June 1999
http://www.cheswick.com/map/index.html
Copyright (C) 1999, Lucent Technologies

# Internet Topology and Power Laws

A few nodes have lots of connections

**Source: Faloutsos et al (1999)**

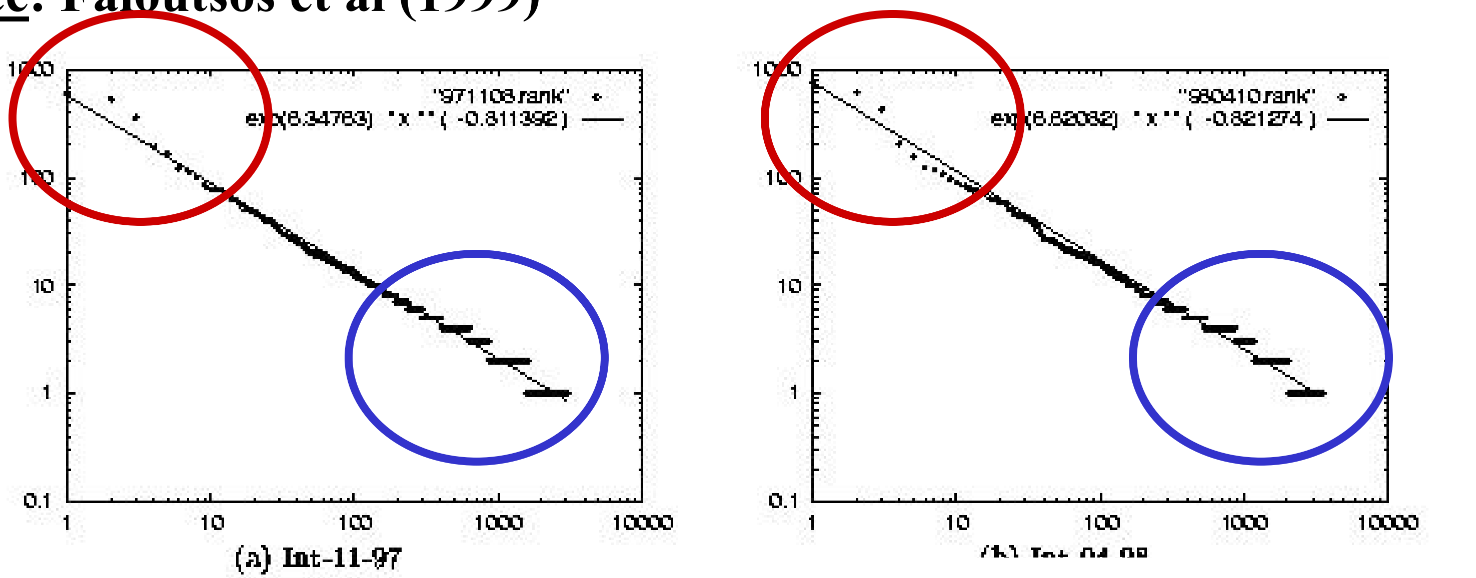


Most nodes have few connections

Figure 3: The rank plots. Log-log plot

- How to account for high variability in node degrees?
- How to explain high variability in node degrees?

# The “Physics” View: Scale-Free Networks (SFN)

- Barabasi and Albert, *“Emergence of scaling in random networks”* (1999)
- Albert, Jeong, and Barabasi, “*Attack and error tolerance of complex networks*” (2000)
- Barabsi, “*Linked: The New Science of Networks*” (2002)

# SFN and Internet Topology

- Ignore Internet-specific “details” for the sake of addressing “fundamental” issues
- High variability in node degrees: Power law!
- Preferential attachment leads to power law degree distributions (“scale-free” structure)
- Central core nodes with high degree (“hubs”): The “Achilles’ heel” of the Internet
- There exist deep links between Internet topology, the power grid, airline network, metabolic networks, etc.


One of
the most-read
papers ever on
the Internet!

# SFN and the Internet

- The high-degree nodes are at the edge and not in the center of the network
- The Internet is extremely robust to hardware failures, but (by design!) very fragile to hijackings by malicious end users
- Bottom line
  - It is conceivable to build a SF Internet
  - No engineer would ever design a SF Internet
  - No user would ever get anything done using a SF Internet
  - Despite getting it completely wrong, SFNs have had impact

# The Networking View: High Variability II

- What causes node degrees to be highly variable?
- Topology is more than connectivity
  - Routers (nodes) have max capacity
  - Links have max bandwidth
- Economic factors impacting network design
  - End user bandwidth demands
  - Link costs
  - Incentive for multiplexing

# 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 |

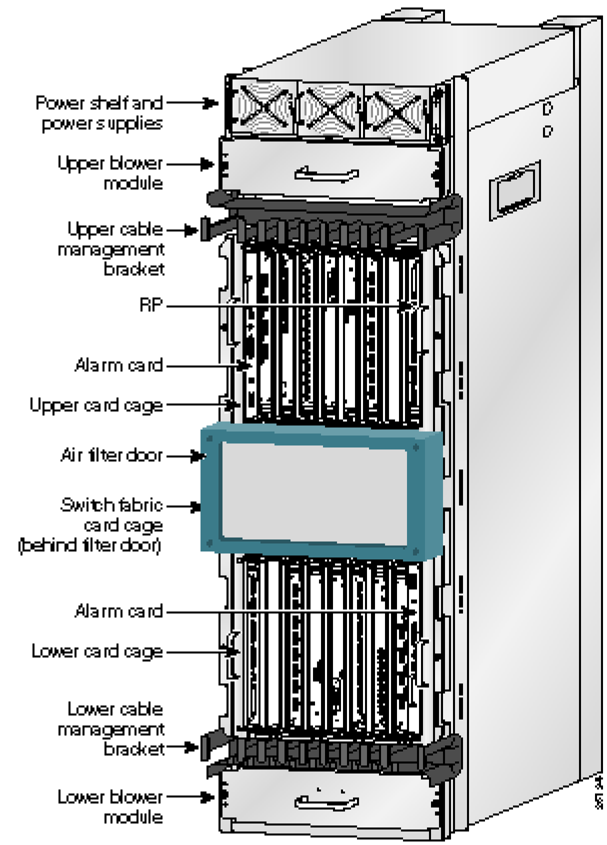


Source: www.cisco.com

# Technologically Feasible Region

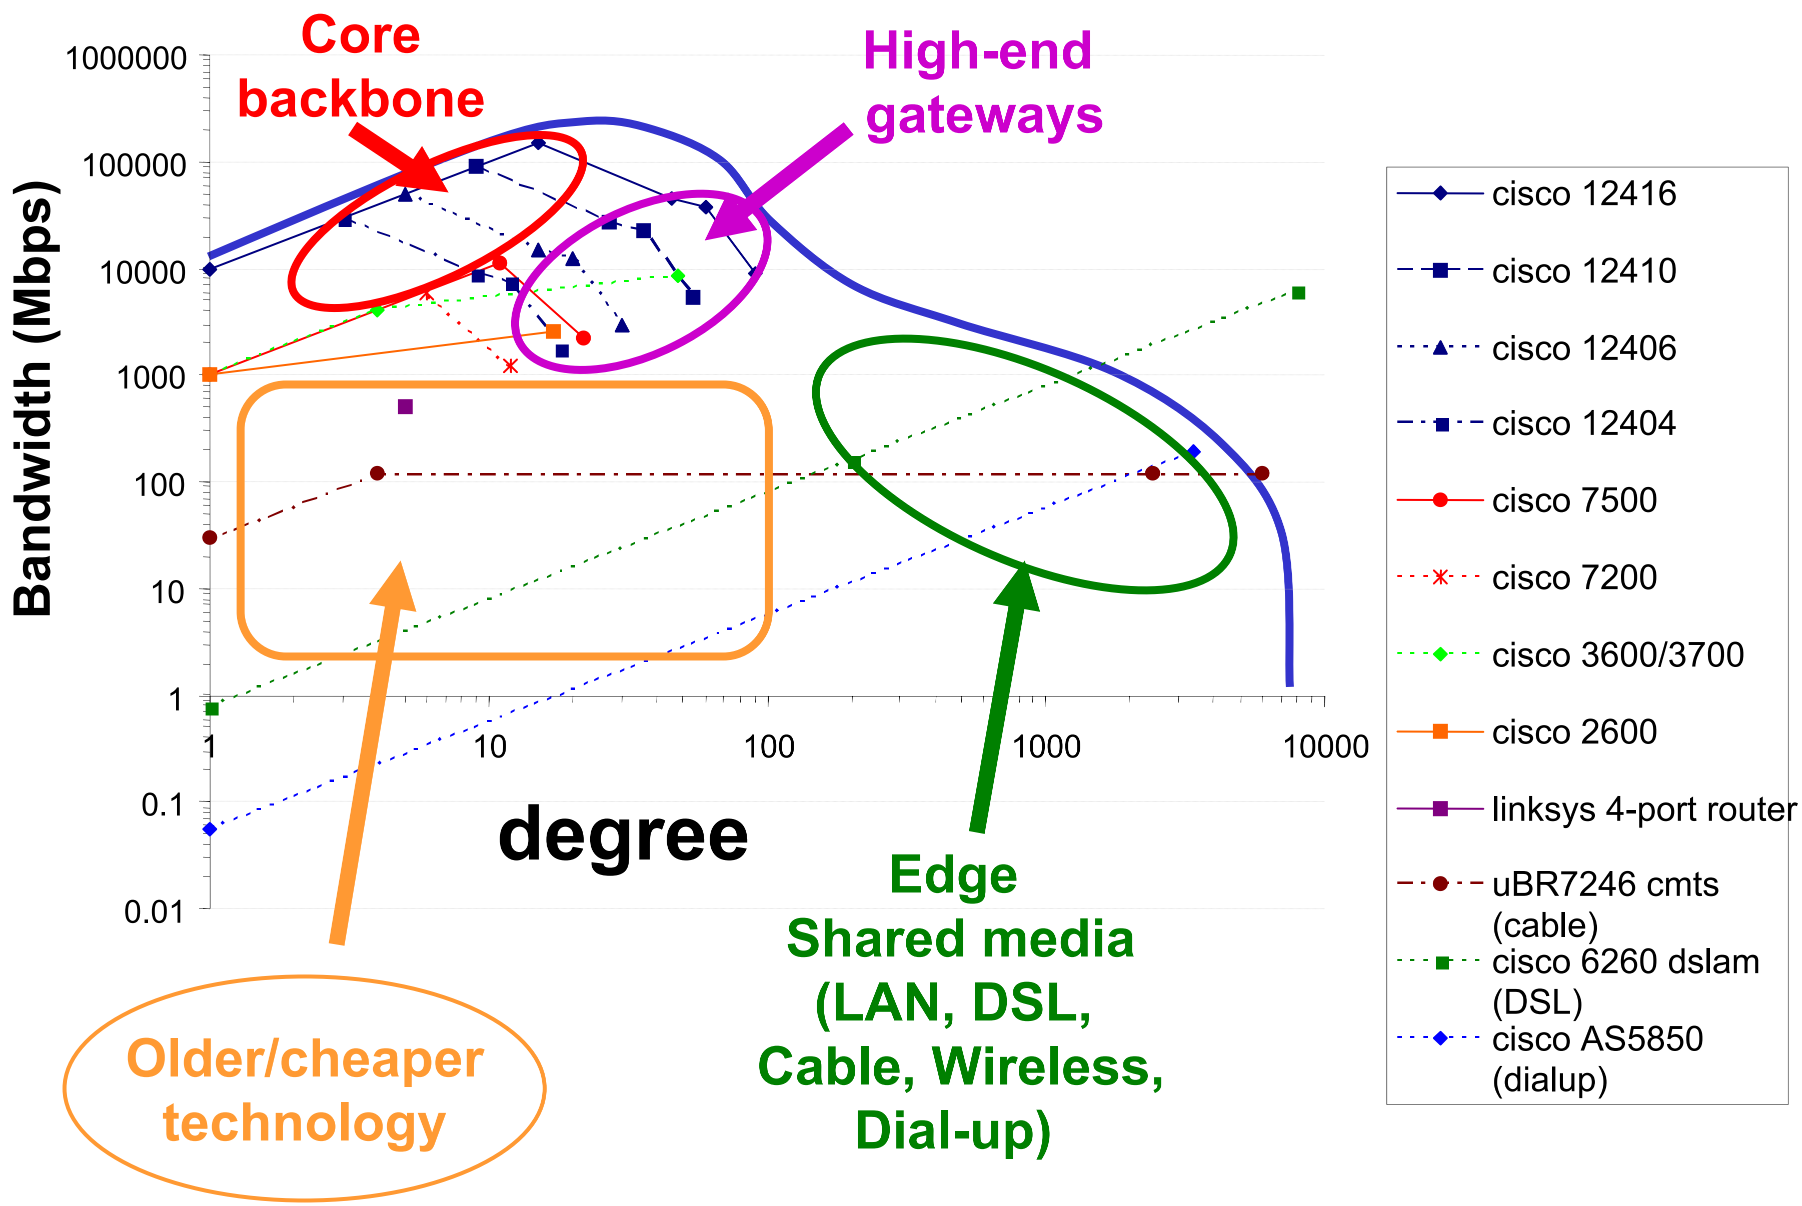

# Internet End-User Bandwidths

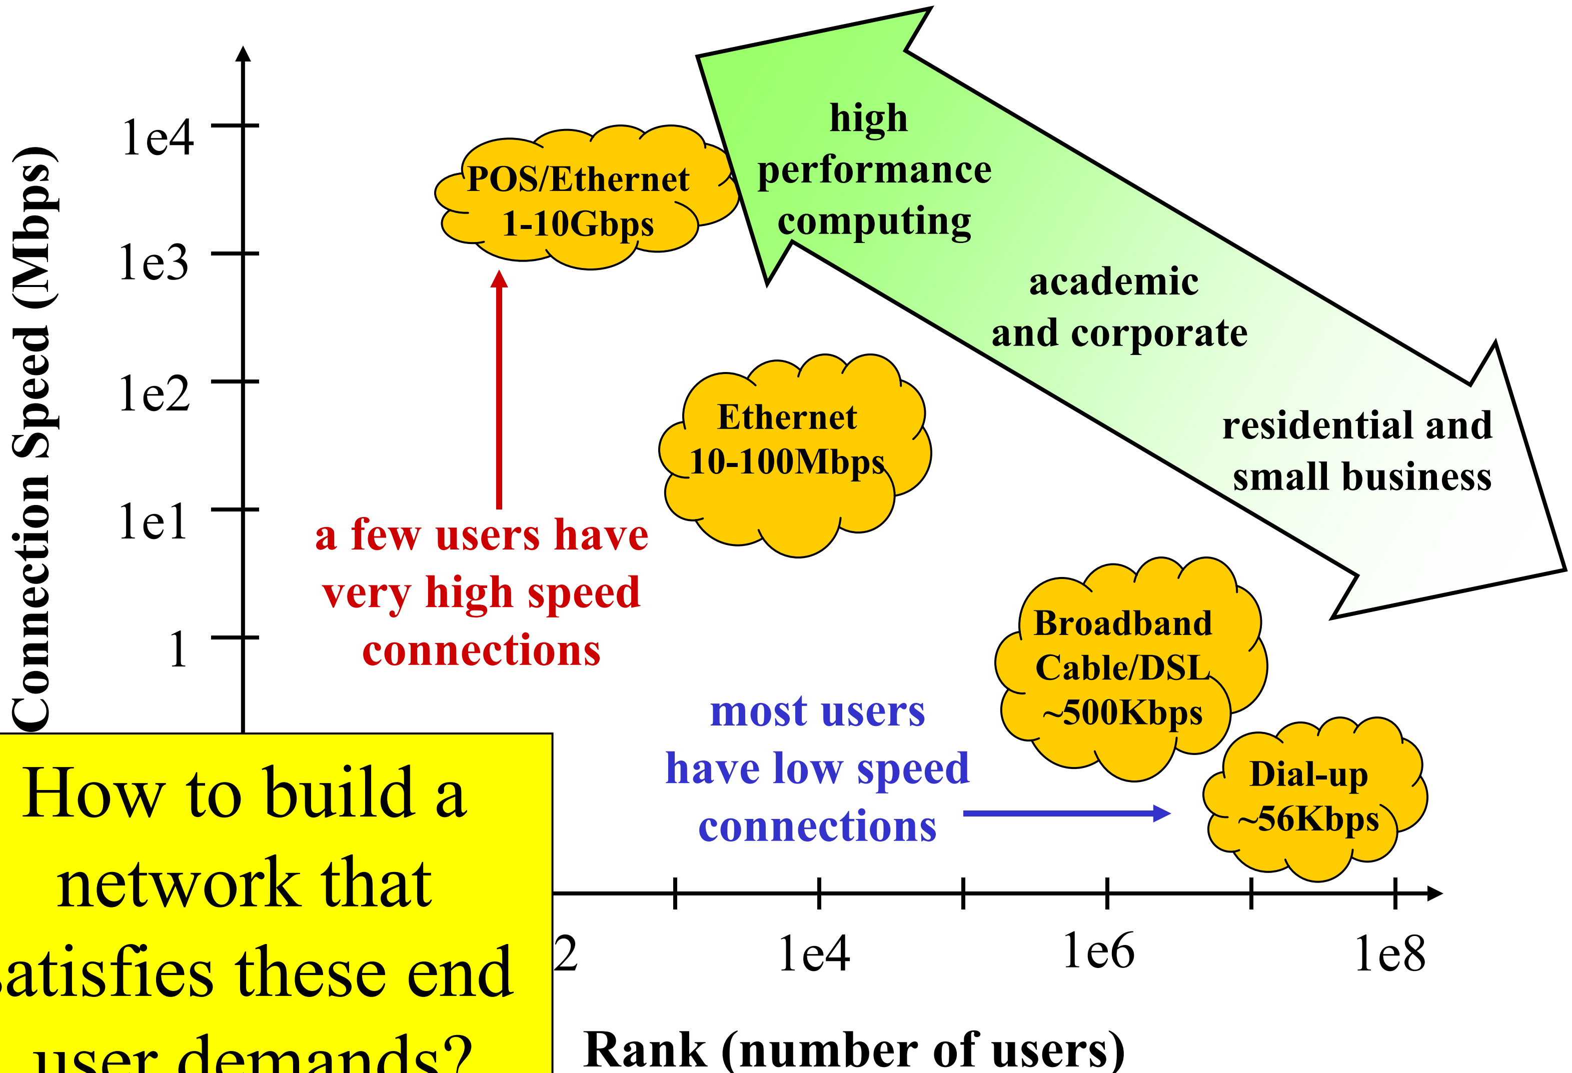


How to build a network that satisfies these end user demands?

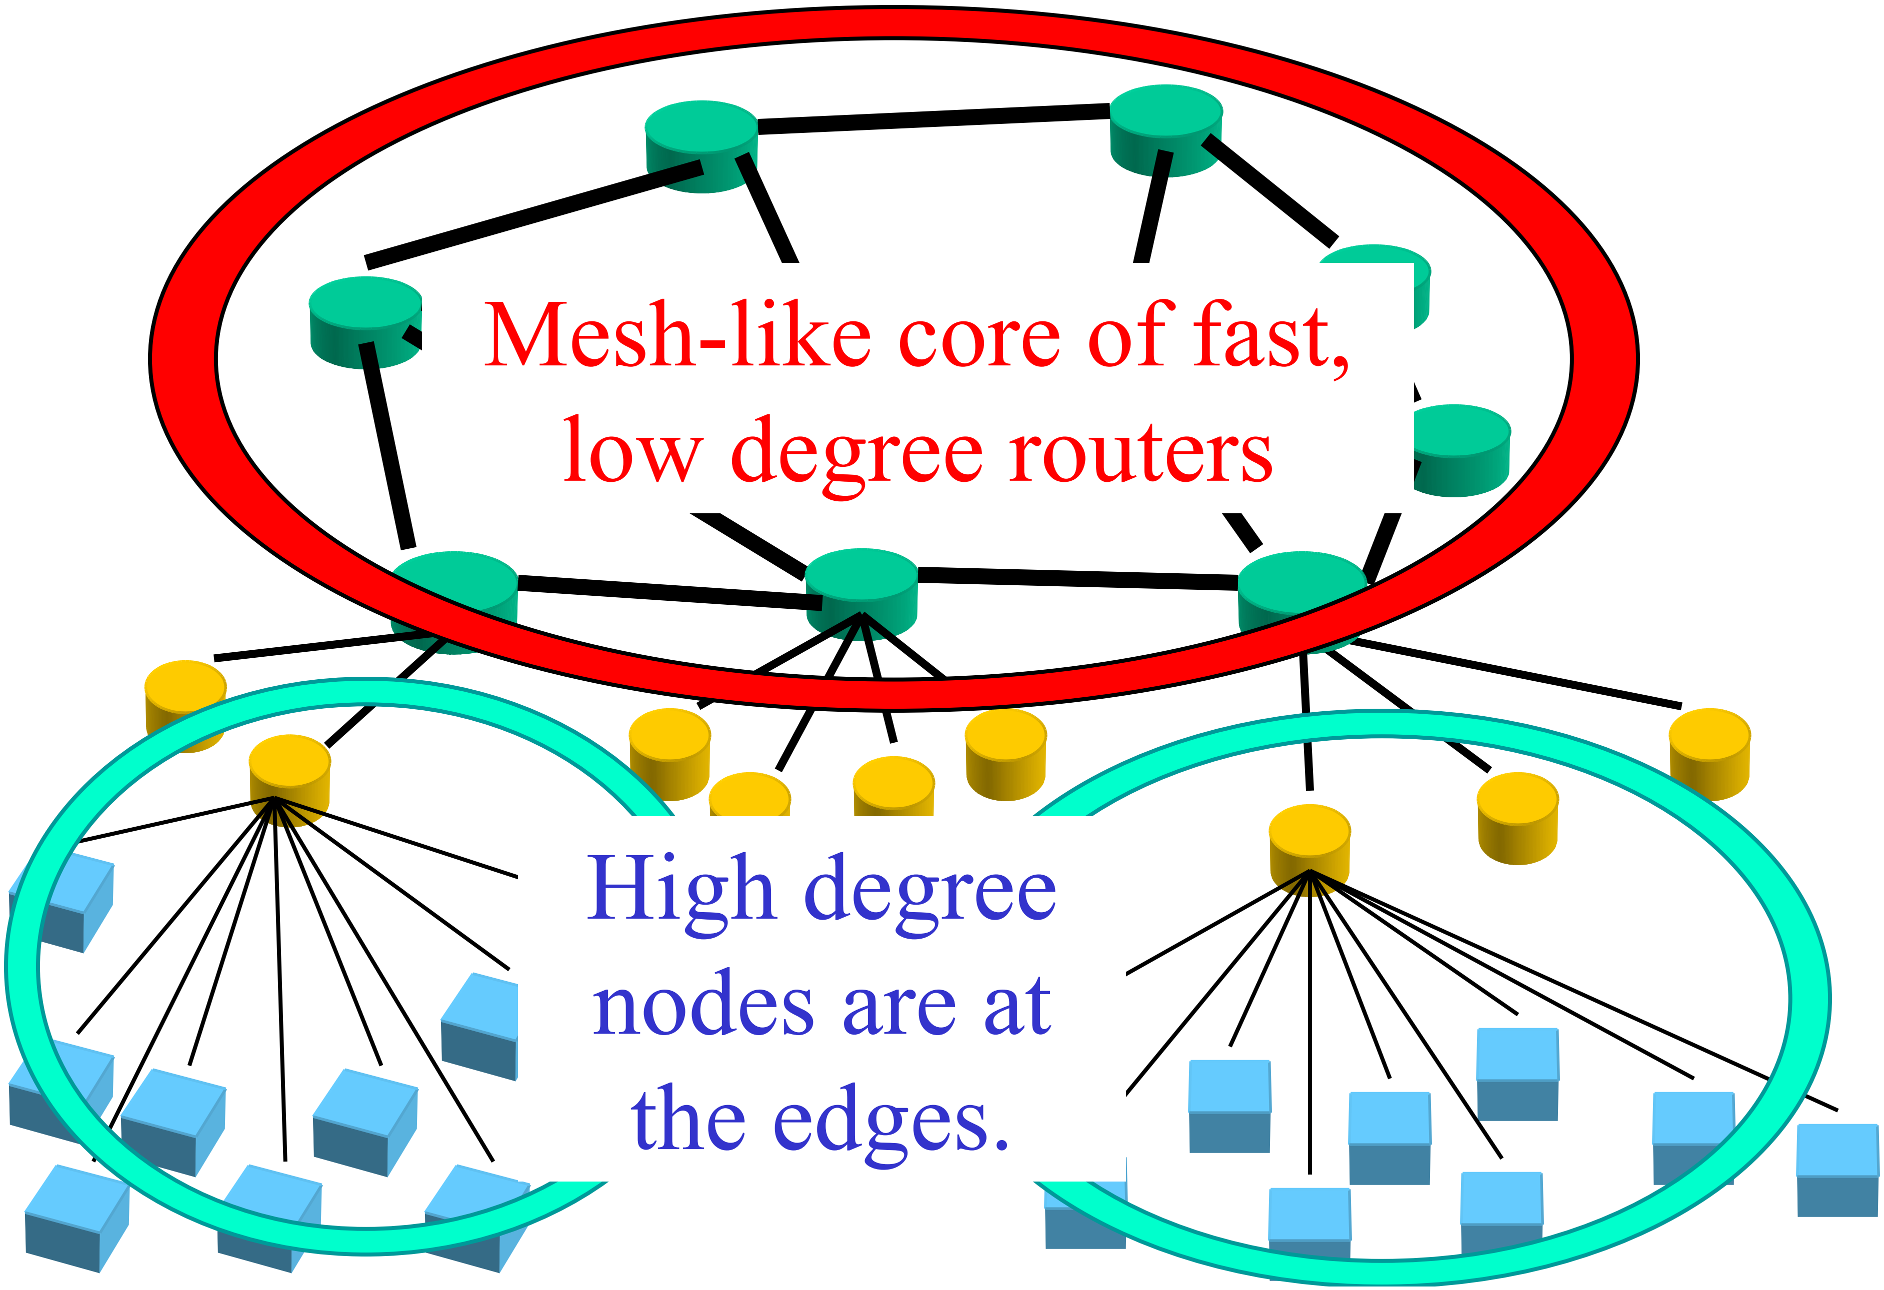
Mesh-like core of fast, low degree routers
High degree nodes are at the edges.

# Heuristically Optimal Topology (HOT)

- Based on HOT framework (Carlson and Doyle)
- Economic  and technology constraints:
  - Mesh-like core of high-speed, low-degree routers
  - High-degree, low-speed nodes at the edge
- Consistent with real observed network

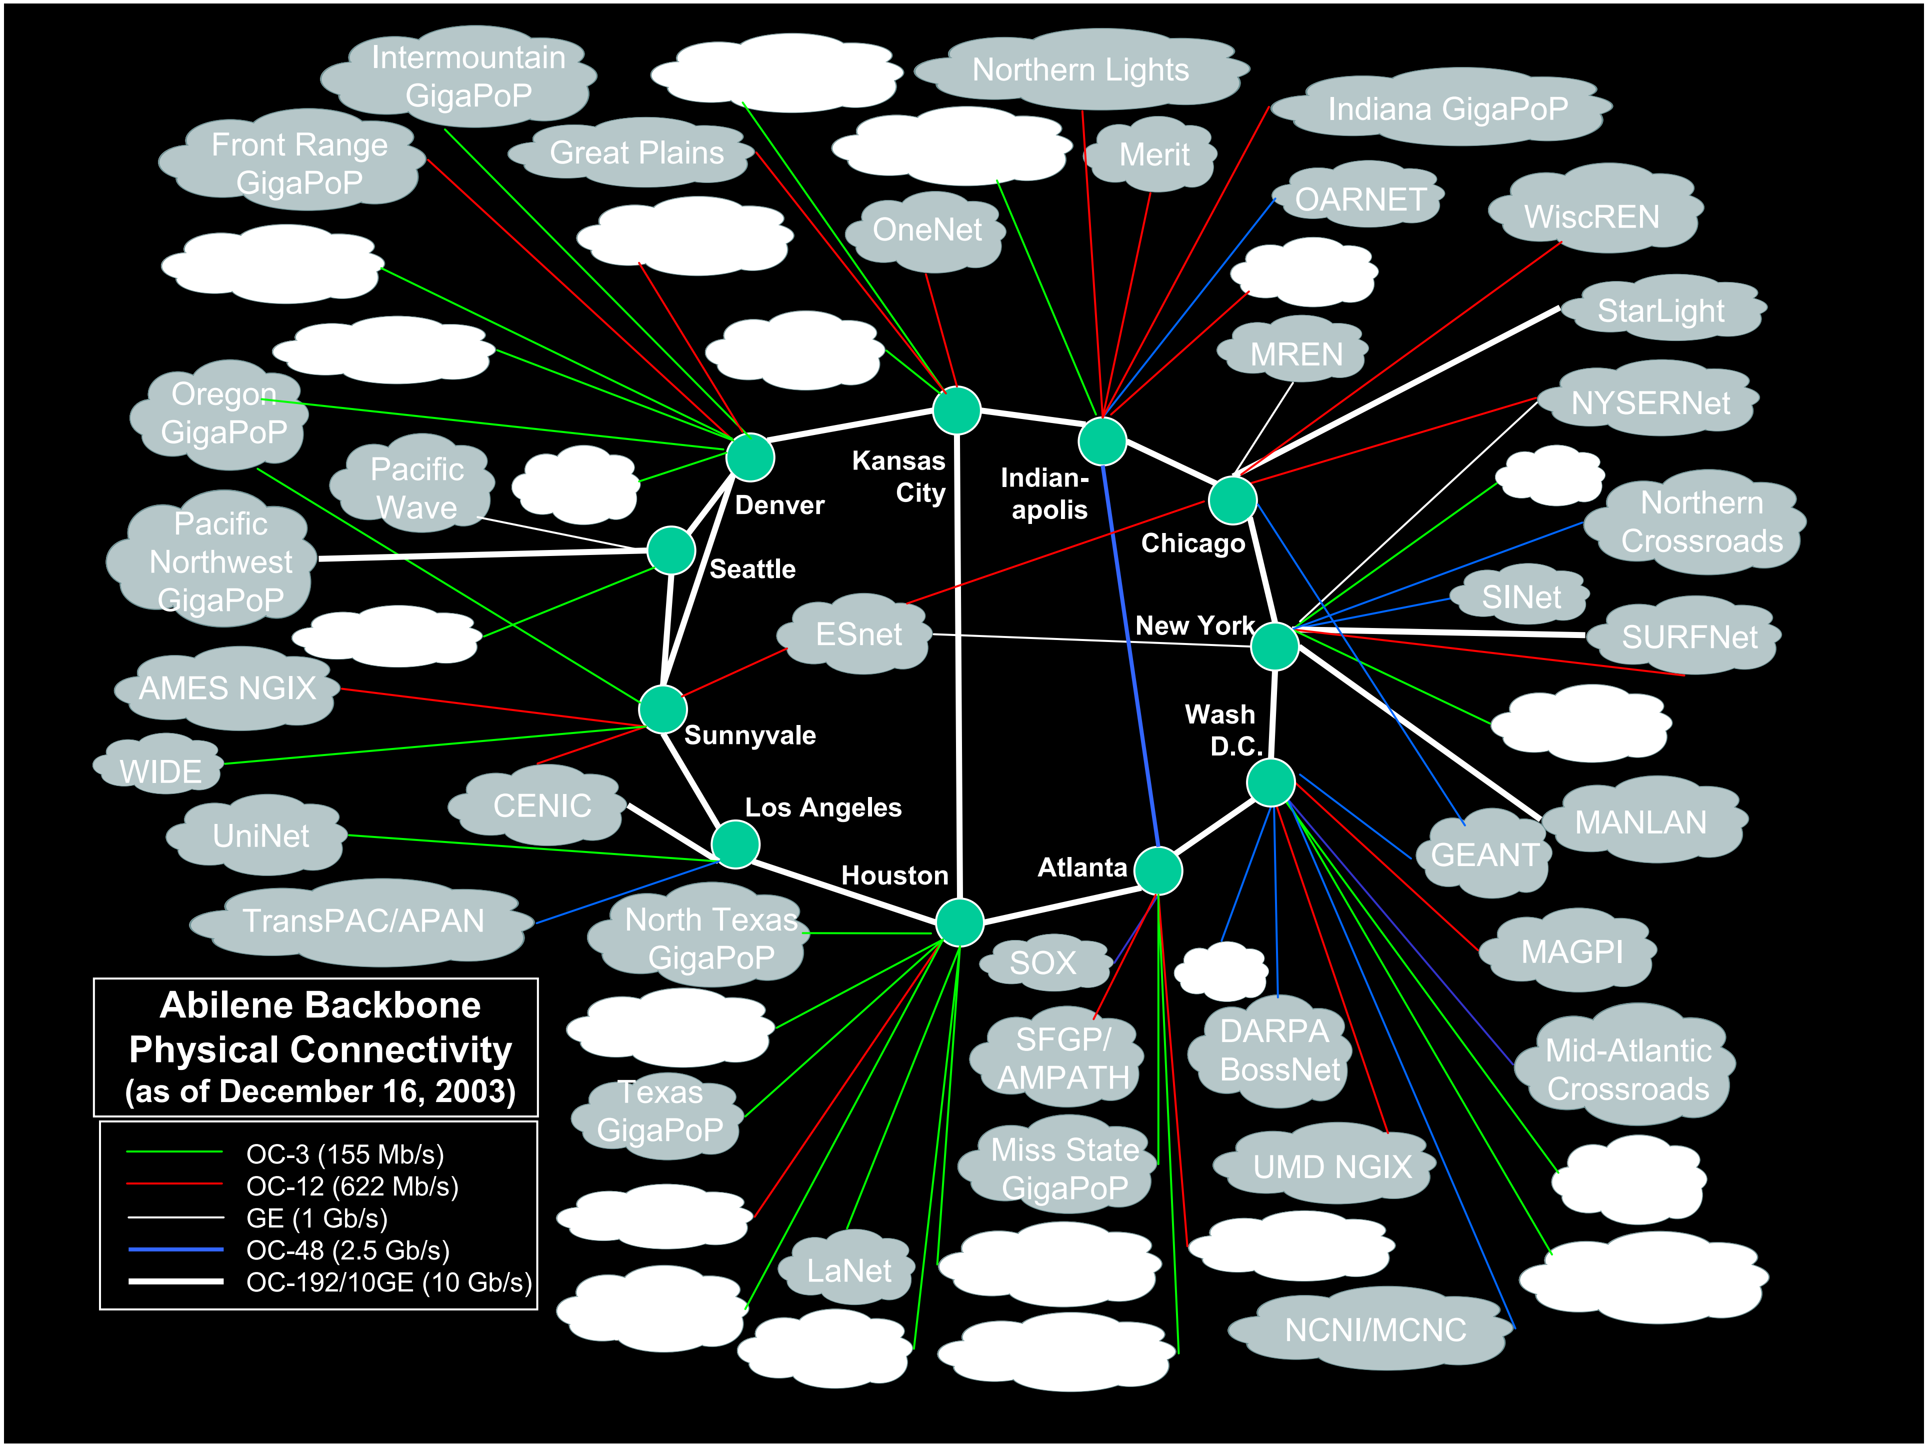

Intermountain GigaPoP
Northern Lights
Indiana GigaPoP
Front Range GigaPoP
Great Plains
Merit
OARNET
WiscREN
OneNet
StarLight
MREN
Oregon GigaPoP
NYSERNet
Pacific Wave
Northern Crossroads
Pacific Northwest GigaPoP
SINet
ESnet
SURFNet
AMES NGIX
WIDE
CENIC
UniNet
GEANT
MANLAN
TransPAC/APAN
North Texas GigaPoP
SOX
MAGPI
SFGP/ AMPATH
DARPA BossNet
Mid-Atlantic Crossroads
Texas GigaPoP
Miss State GigaPoP
UMD NGIX
LaNet
NCNI/MCNC
Denver
Kansas City
Indian- apolis
Chicago
Seattle
New York
Sunnyvale
Wash D.C.
Los Angeles
Houston
Atlanta
Abilene Backbone Physical Connectivity (as of December 16, 2003)
OC-3 (155 Mb/s)
OC-12 (622 Mb/s)
GE (1 Gb/s)
OC-48 (2.5 Gb/s)
OC-192/10GE (10 Gb/s)

# Internet Topologies

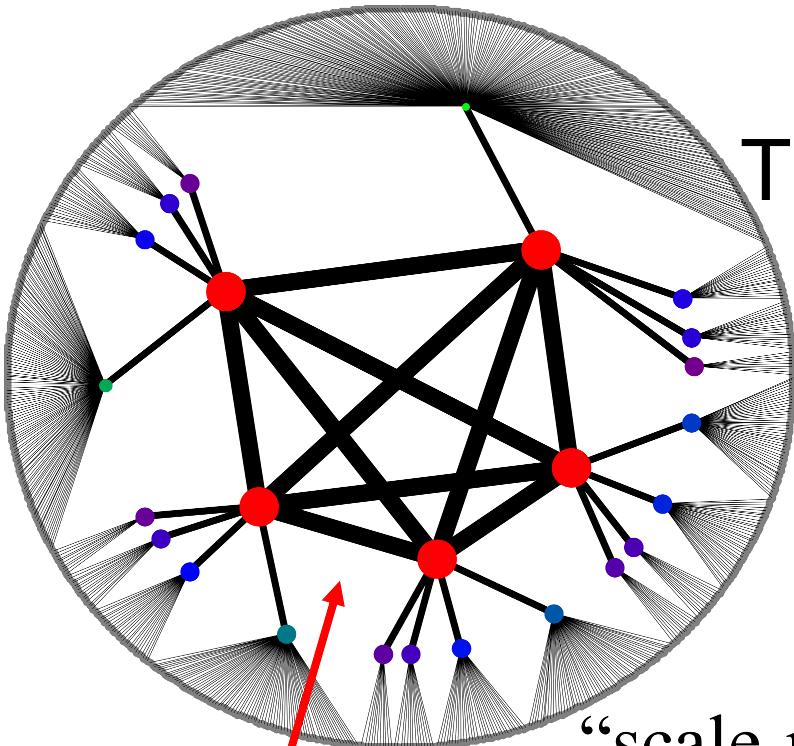

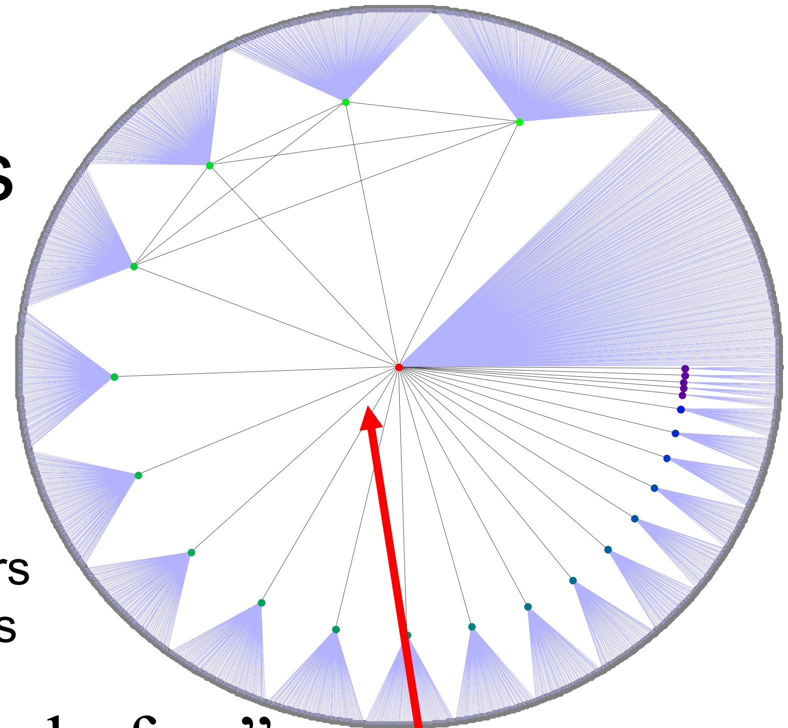

nodes=routers
edges=links

25 interior routers
818 end systems

“scale-rich” vs. “scale-free”

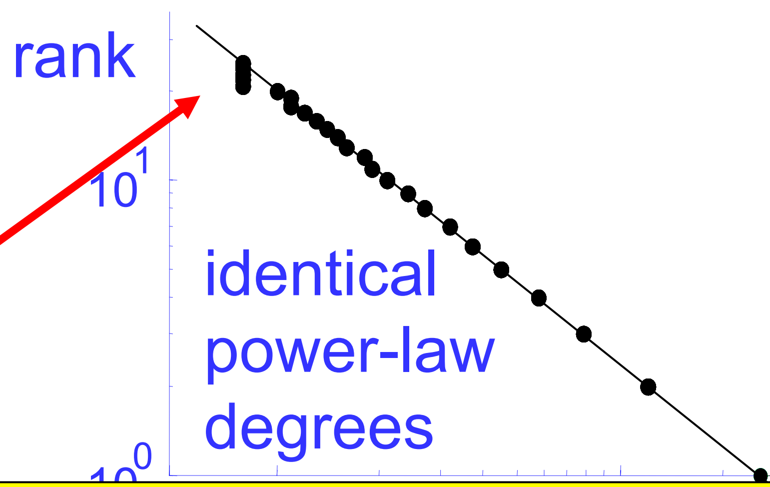


Low degree mesh-like core

High degree hub-like core

**How to characterize / compare these two networks?**

# Characterization of Topologies

- Network performance
  - Technology constraint
  - End user demands, traffic matrix
  - How well is the network able to carry a given traffic?
- Network likelihood
  - All simple connected graphs with same node degree distribution
  - Assign probability to each graph
  - Related to “degree of scale-free” and “degree of self-similarity”

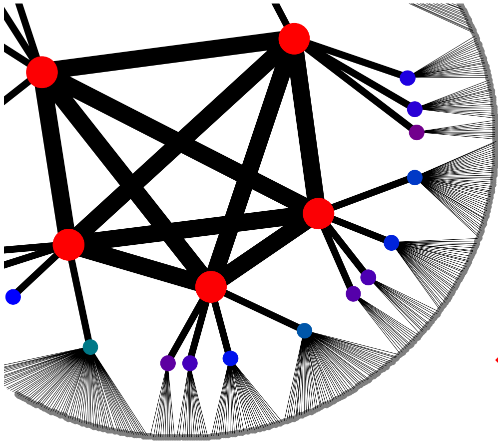

Why such striking differences with same node degree distribution?

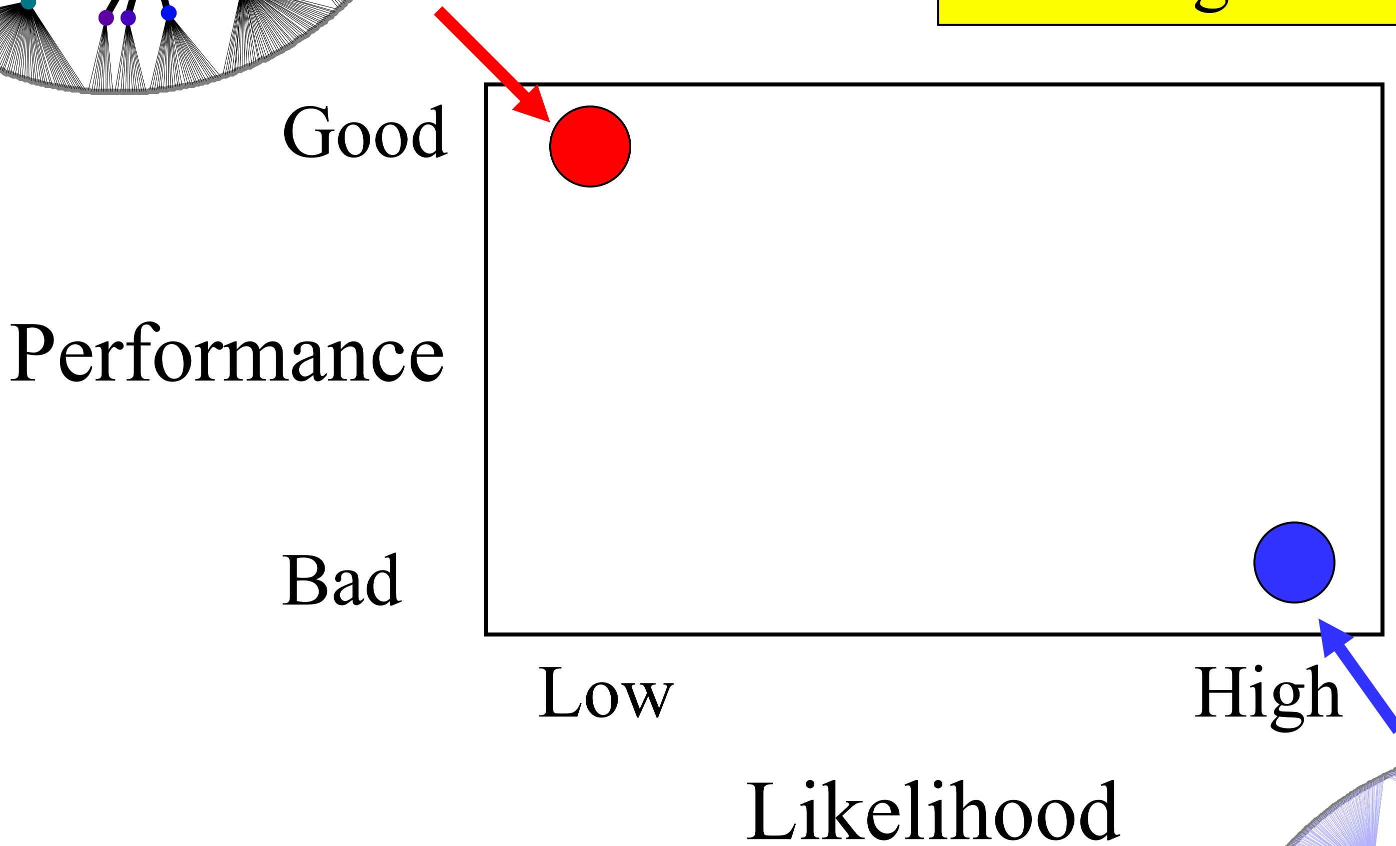


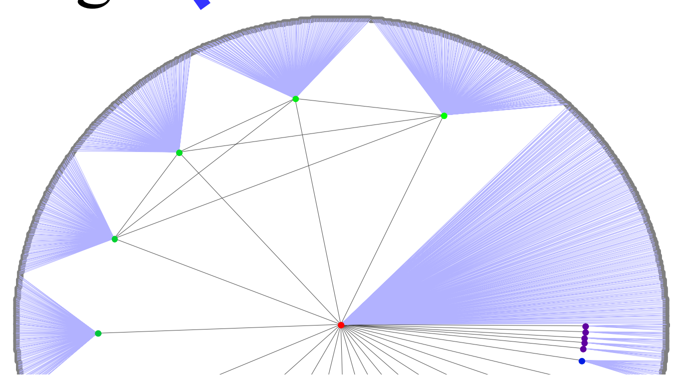

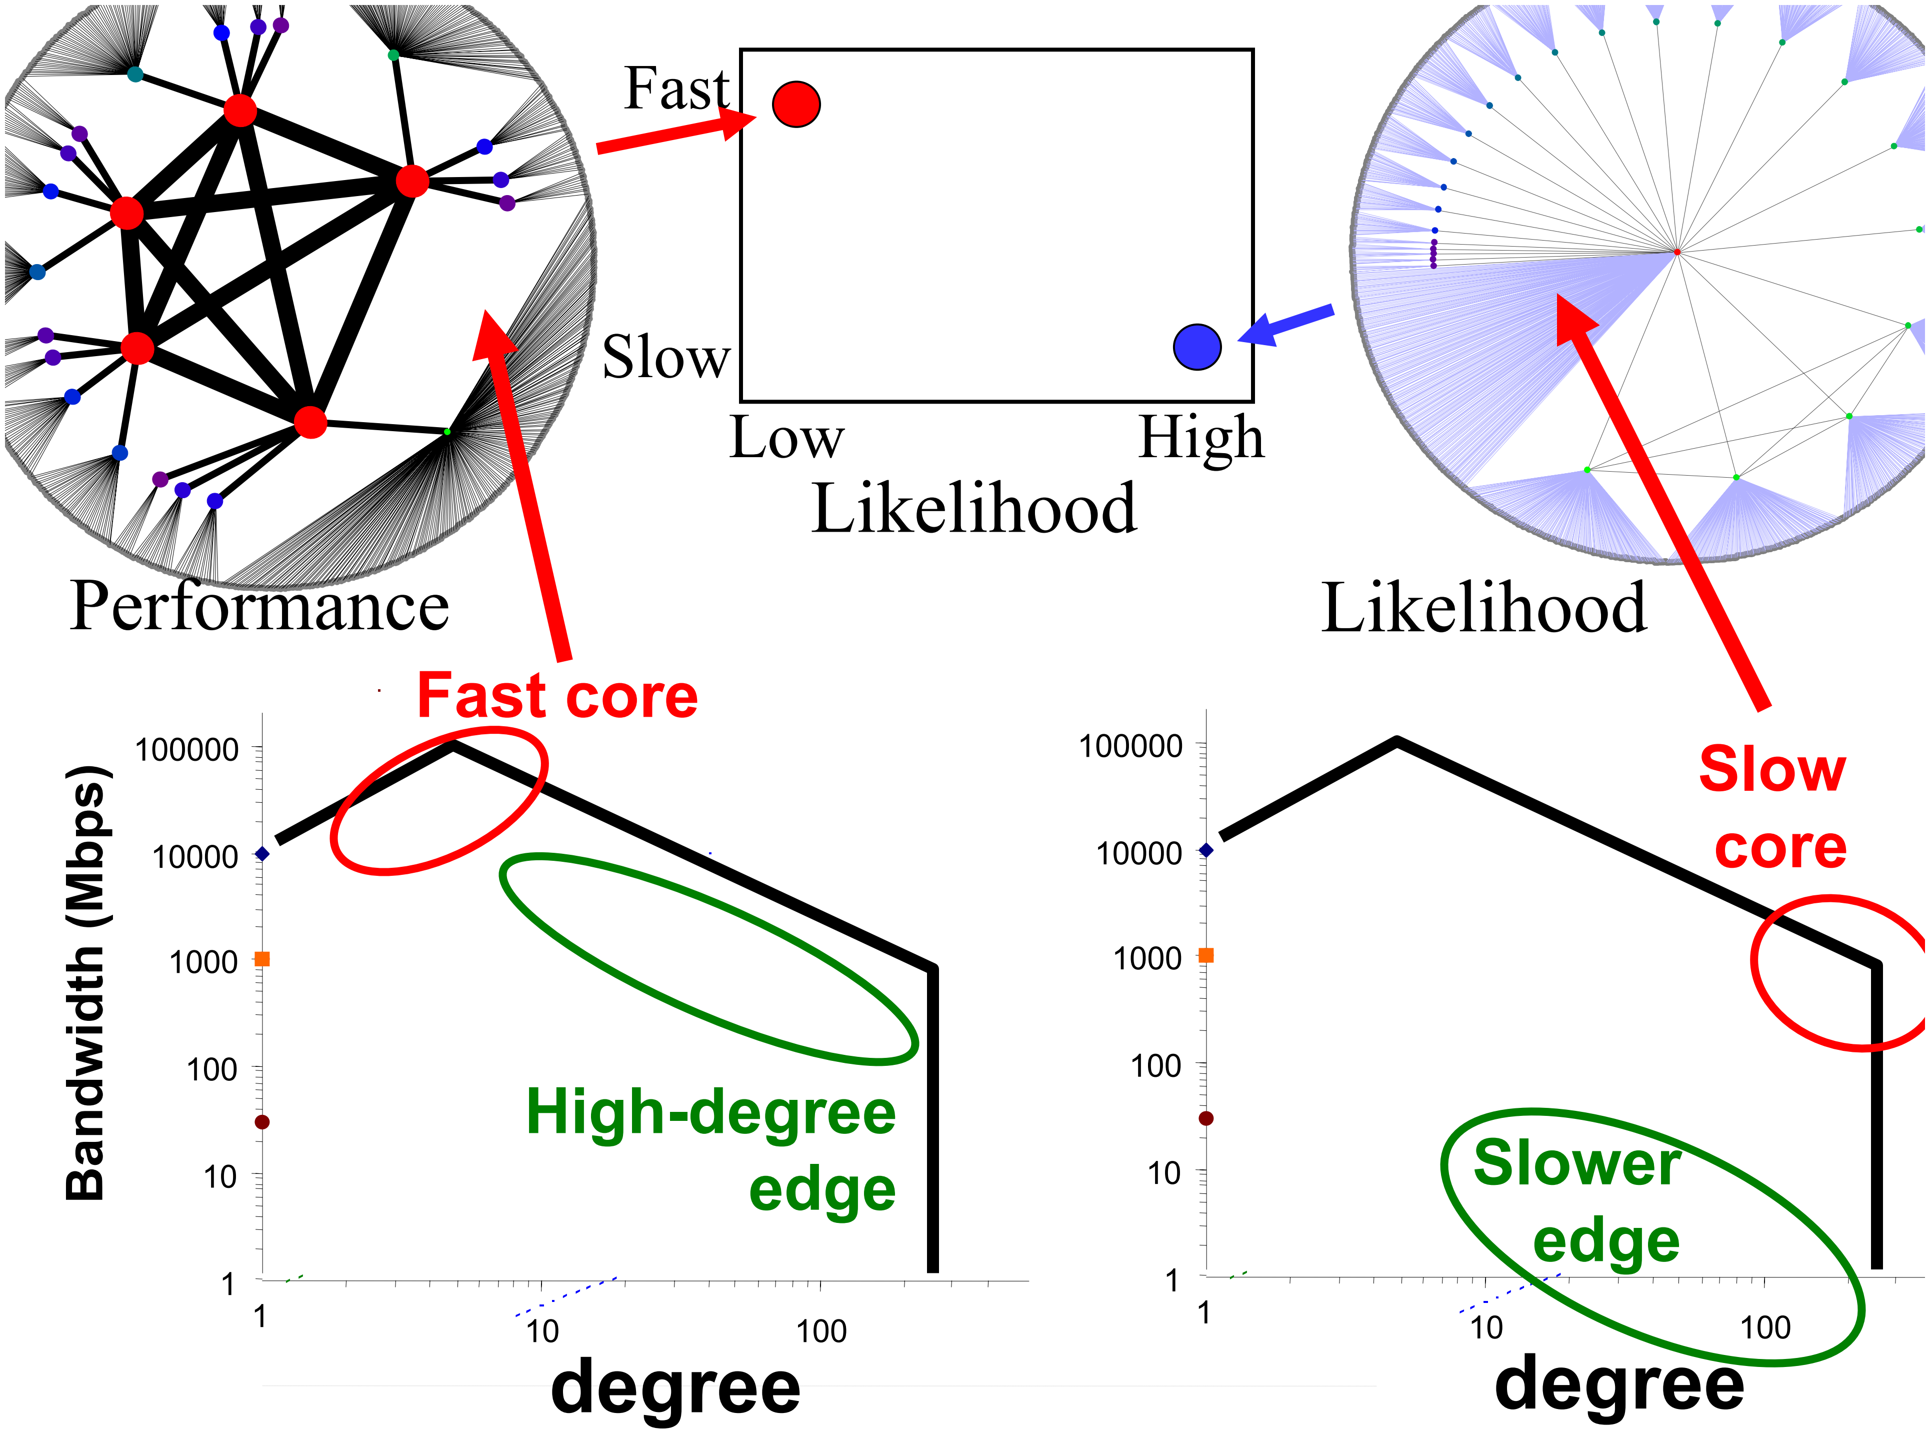

Fast
Slow
Low
High
Likelihood
Performance
Likelihood
Fast core
Slow core
Bandwidth (Mbps)
100000
10000
1000
100
10
1
1
10
100
degree
High-degree edge
Slower edge
100000
10000
1000
100
10
1
1
10
100
degree

# High Variability and Internet Topology

- Newer measurements show that router node degrees are not consistent with a power law distribution
- Domain knowledge is crucial
- High variability in end user demands leads to high variability in node degrees
- HOT networks are unlikely to arise at random
- Basic ingredient for studying topologies at higher layers (AS-graph, overlay networks)

# 2004: JDC’s 1st SIGCOM Paper

Li, Alderson, Willinger, and Doyle, “A first-principles approach to understanding the Internet’s router-level topology”, ACM SIGCOMM 2004

- Provides an explanation of the origin of the high variability in node degrees
- Resolves much of the confusion and controversy created by SFN
- Motivates the development of a rigorous mathematical treatment of SFN
- SFN may be ok for “virtual” networks

# A Look Ahead

- Internet-specific multiscale analysis
  - Internet architecture defines an (unusual) multiscale structure
  - Bytes, packets, IP-flows, OD-flows, AS-flows
  - Link, individual IP addresses, prefixes, ASes
  - Separate pieces in place, but no coherent approach exists
- Towards a theory of the Internet
  - First theoretical results on TCP confirm the ingenuity of the original architects of the Internet
  - What will be the first negative result? Routing?