Why Understanding Anything About The Internet Is Painfully Hard

Vern Paxson

AT&T Center for Internet Research at ICSI
International Computer Science Institute
Berkeley, CA

vern@aciri.org

April 28, 1999

Overview:

What is “the Internet”?

An immense moving target.

The desperate search for invariants.

The collapse of traditional mathematical modeling of networks.

Fractal traffic & Pareto tails.

What is the Internet?

A collection of computers capable of communicating with each other using a standard set of “protocols.”

30.1 million “hosts” as of January 15, 1998.

An <u>inter</u>network, made up of 50,000+ separate networks.
A <u>wide-area</u> network $\Rightarrow$ speed of light matters.

Each network comprises numerous hosts and (fewer) “routers.”
Hosts are endpoints, routers are internal way-stations.
Connections between them are “links.”

What the Internet isn't:

In connection-oriented networks, the fundamental building block is the notion of a connection between two computers:

- Resources are reserved when a connection starts
- Resources are reclaimed when a connection finishes

Example: the telephone system.

Advantages:

no surprises (except connection denied); easy billing.

Disadvantages:

wasted resources; lots of “state” inside the network.

What the Internet is:

In packet networks, the fundamental building block is a single packet of data:

- Each packet has full “addressing” information.
- Once a router forwards a packet, *it can forget about it*.
- Delivery today is generally best effort.

Example: the postal system.

Advantages: no state = very resilient; better use of resources; easy interconnection.

Disadvantages: packet loss; billing is an open problem.

Packet loss and congestion control:

What if more packets arrive at a router than it can send ahead along its links?

Answer: packets are queued in memory buffers.
If the router runs out of memory, it drops the packet.

The Internet protocols *must* use this feedback in order to adapt their sending rate to current conditions along a network path, or else risk congestion collapse.

$\Rightarrow$ Traffic is constantly adapting, reacting & changing.
$\Rightarrow$ Trace-driven simulation is inherently tricky.

USENET Bulletin Board Traffic Volume

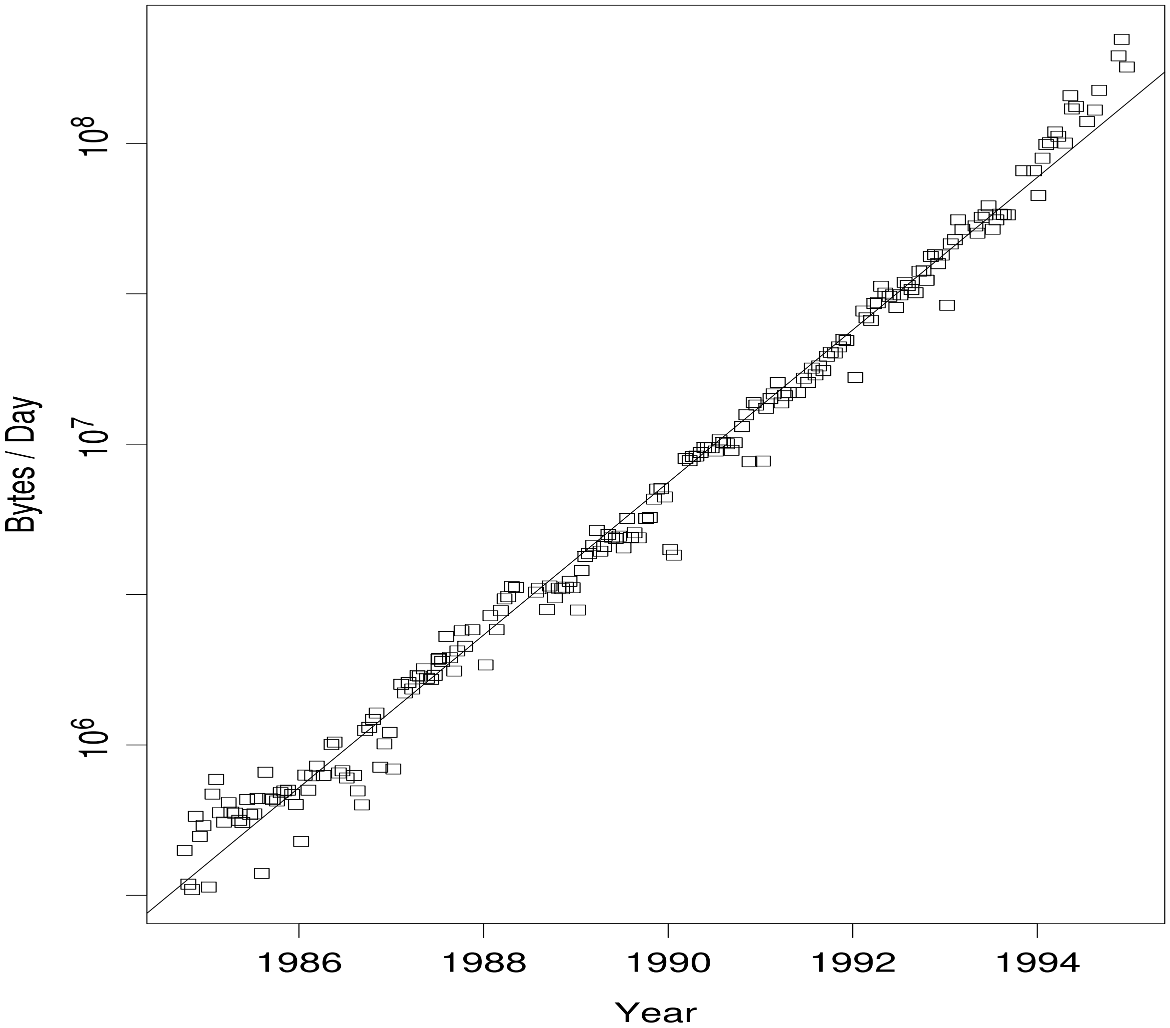

Courtesy of Rick Adams.

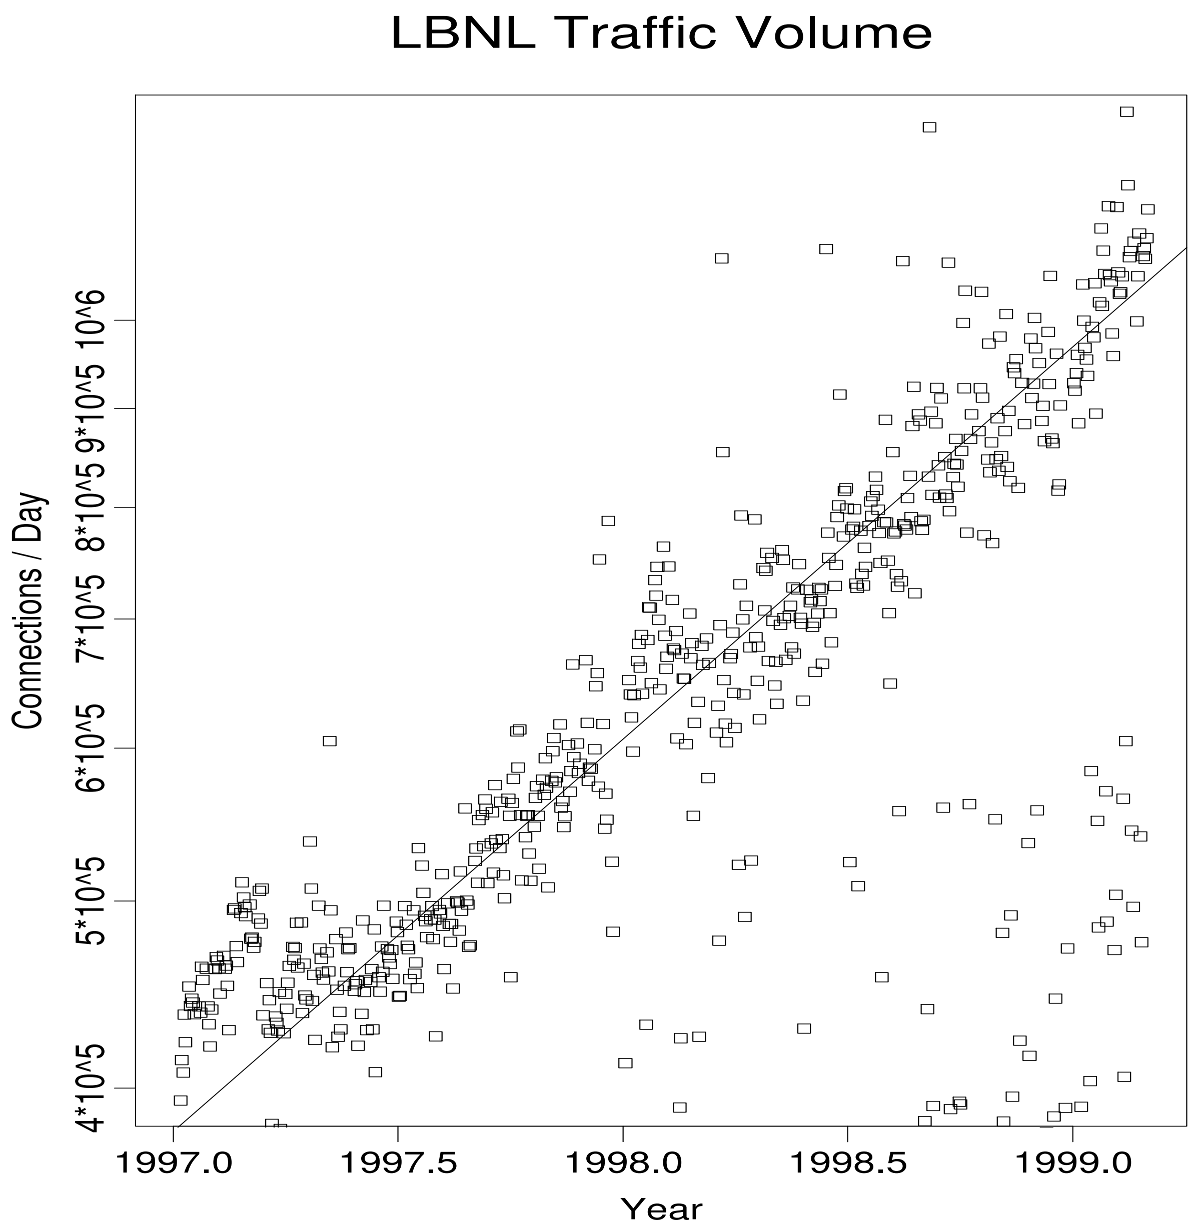

LBNL Traffic Volume
Connections / Day
4*10^5
5*10^5
6*10^5
7*10^5
8*10^5
9*10^5
10^6
1997.0
1997.5
1998.0
1998.5
1999.0
Year

The Internet — huge & diverse:

In January, 1990, the Internet had about 200,000 hosts.

Five years later, it had about 4,900,000.

$\Rightarrow$ If you blink, you're out of date.

What about since then?

9.5 million in 1996. 16.1 million in 1997. 30.1 million in 1998.

55.2 million as of April 26, 1999.

And: good evidence that current estimation techniques underestimate, perhaps by 20%.

The Internet — huge & diverse, con't:

In January, 1991, about 15% of the data bytes going into LBNL were USENET Network News Bulletin Board.
At USC, it was 60%.

On an MCI backbone link in Sep. 1997, 70% of the packets were HTTP, the dominant Web protocol.
For UCB's Internet link in Dec. 1997, it was 37%.

$\Rightarrow$ There's no such thing as a typical site.

The Internet — huge & diverse, con't:

In Oct. 1992, the median FTP data transfer at LBNL, computed over 60,000 samples, was 4,500 bytes.
Five months later, w/ 80,000 samples, it was 2,100 bytes.

In Mar. 1998, median transfer at LBNL was 10,900 bytes.
In Dec. 1998, it was 5,600 bytes (329,000 samples).

$\Rightarrow$ There's no such thing as a typical time period.

Growth of LBNL's WWW Traffic

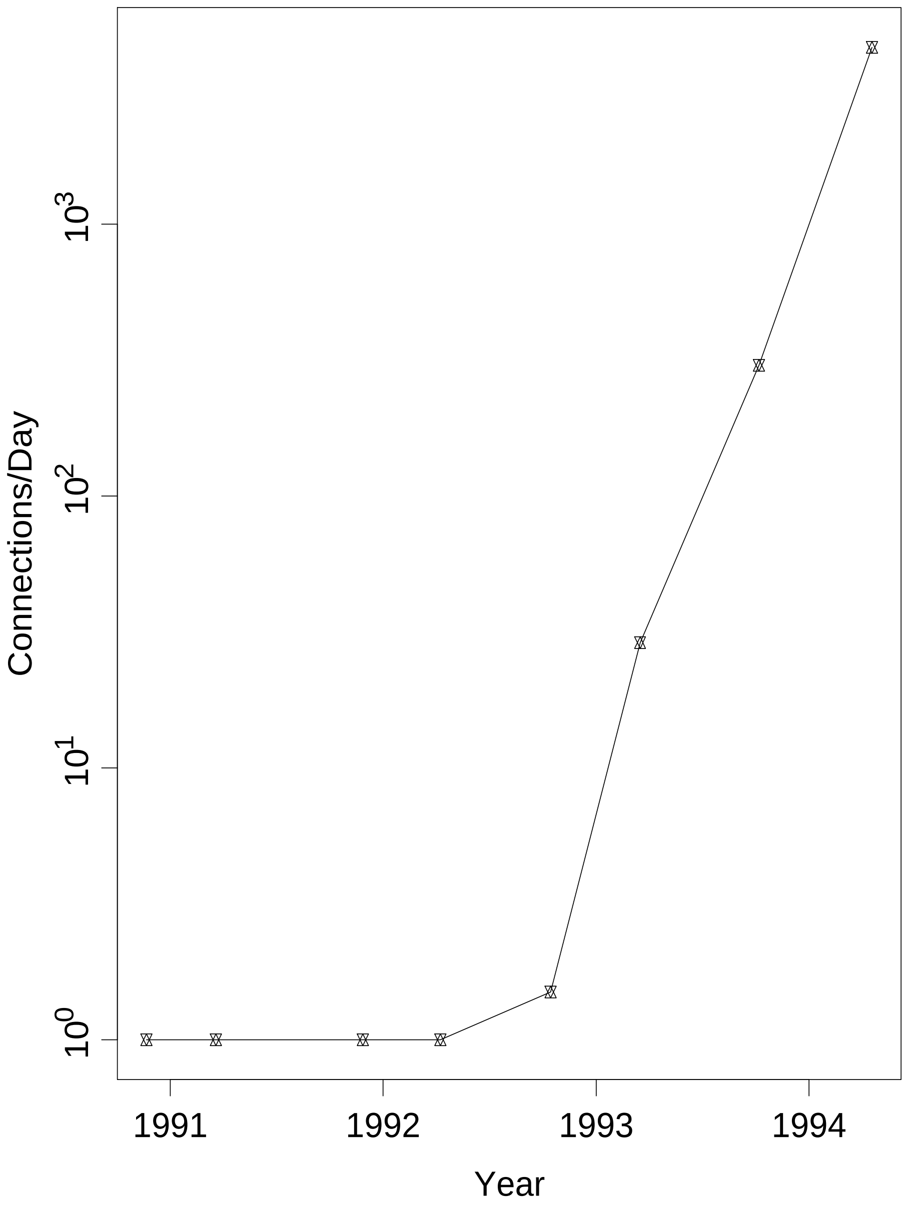

New applications radically alter the landscape:

In the *month* of October, 1992, LBNL made *45* World-Wide Web connections. $6\frac{1}{2}$ years later, it makes about *900,000* each *day*.

In January, 1992, the “MBone” multicast backbone didn't exist. In January, 1995, it accounted for 20% of the traffic at DEC-WRL, 40% at LBNL, and more than 50% at Bellcore.

(In December, 1997, it was about 12% at LBNL.)

What else is coming?

Telephony: whole phone system is migrating to the Internet.

Pricing: but in what form?

Multicast: ubiquitious and native?

Quality of service: support for *differentiated services.*

Large-scale Web caching & replication a very active area of research.

USENET Traffic Volume

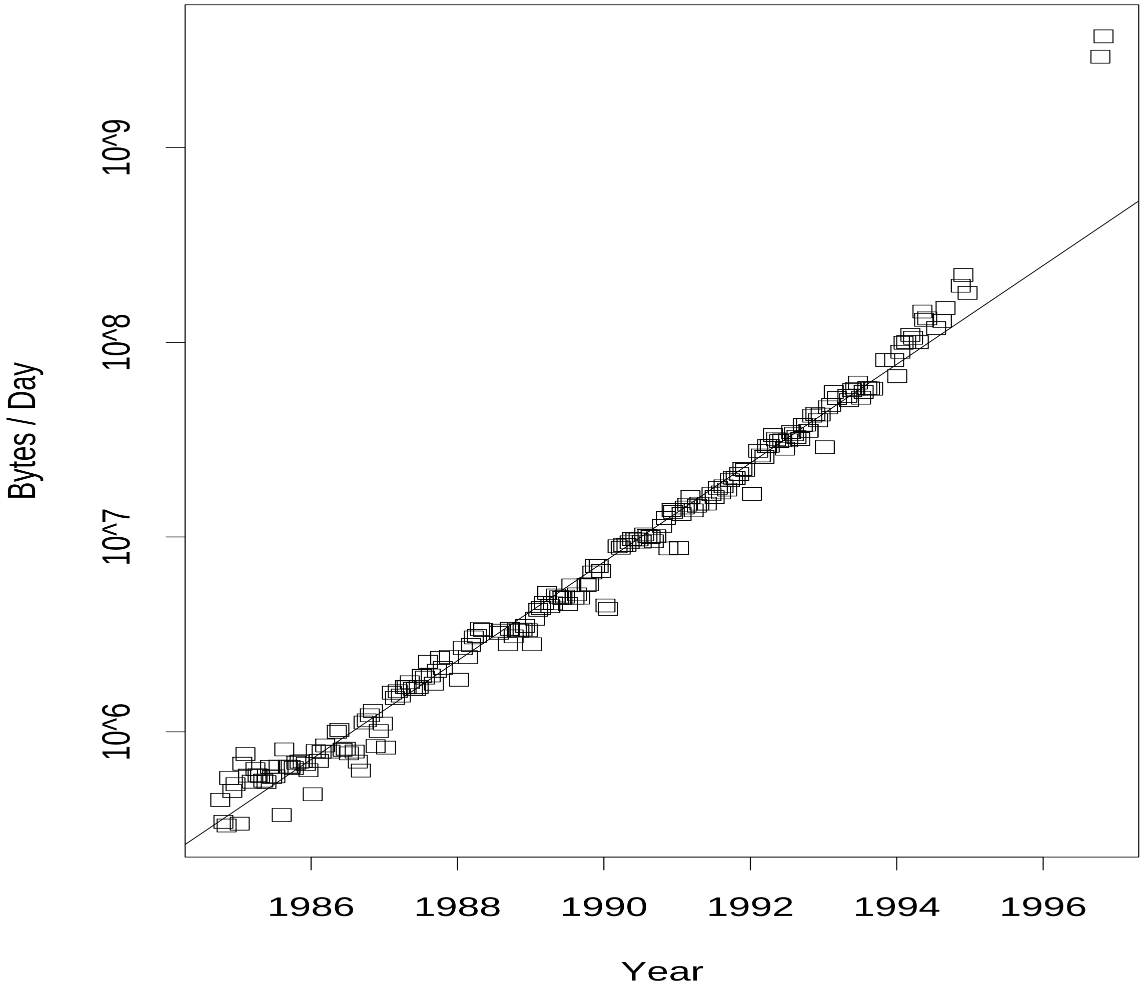

Data courtesy of Rick Adams and David Lawrence.

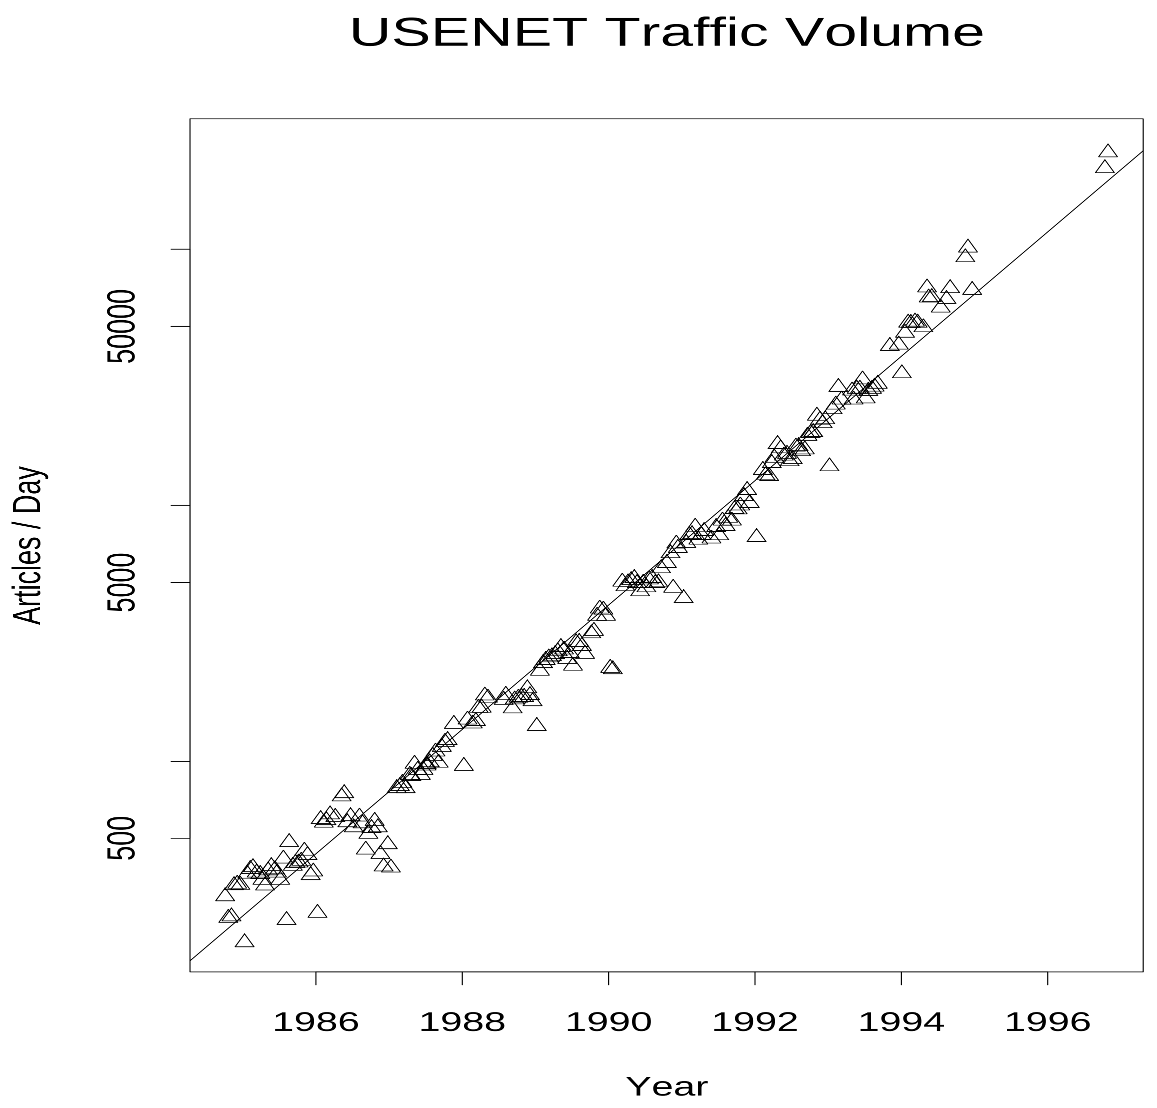

Data courtesy of Rick Adams and David Lawrence.

The lay of the land:

Ubiquitous diversity and change: over time, across sites, how the network is used, by whom.

Thinking in terms of “typical” a lethal pitfall.
But how else do we develop intuition and understanding?

Desperately need <u>invariants</u>: things not changing in a sea of change.

Implications for simulating Internet traffic:

"If you run a single simulation, and produce a single set of numbers (e.g., throughput, delay, loss), and think that that single set of numbers shows that your algorithm is a good one, then you haven't a clue."
$\Rightarrow$ Judiciously explore the parameter space.

Hold all parameters fixed and explore one to gauge sensitivity.
Rule of thumb: work through several orders of magnitude.
Beware: sometimes small changes have large effects (non-linearities).

Examine <u>why</u> effects change with particular parameters.

Implications for modeling Internet traffic:

Desperately need parsimonious models . . .

. . . because we don't have a clue how to meaningfully set more than a few parameters.

The main tool available over last 20 years: Poisson modeling.

Central idea: network events (packet arrivals, connection arrivals) are independent. There's just a single **rate** parameter.

Implications of Poisson modeling:

Times between “calls” distributed exponentially:
$F(x) = 1 - e^{-\lambda x}$, and *independently*.

Number of calls has Poisson distribution:

$$P\{N(t+s) - N(s) = n\} = e^{-\lambda t}\frac{(\lambda t)^n}{n!}.$$

Variations: introduce correlations between events that, asymptotically, decay exponentially.

If assumptions correct, can analyze intricate systems.

Poisson modeling tells us:

Correlations are fleeting.

Modest buffer increases give large gains in capacity.

Traffic multiplexes extremely well:

$$\mathsf{Var}(X^{(m)}) \sim m^{-1}.$$

It's easy to define a “burst.”

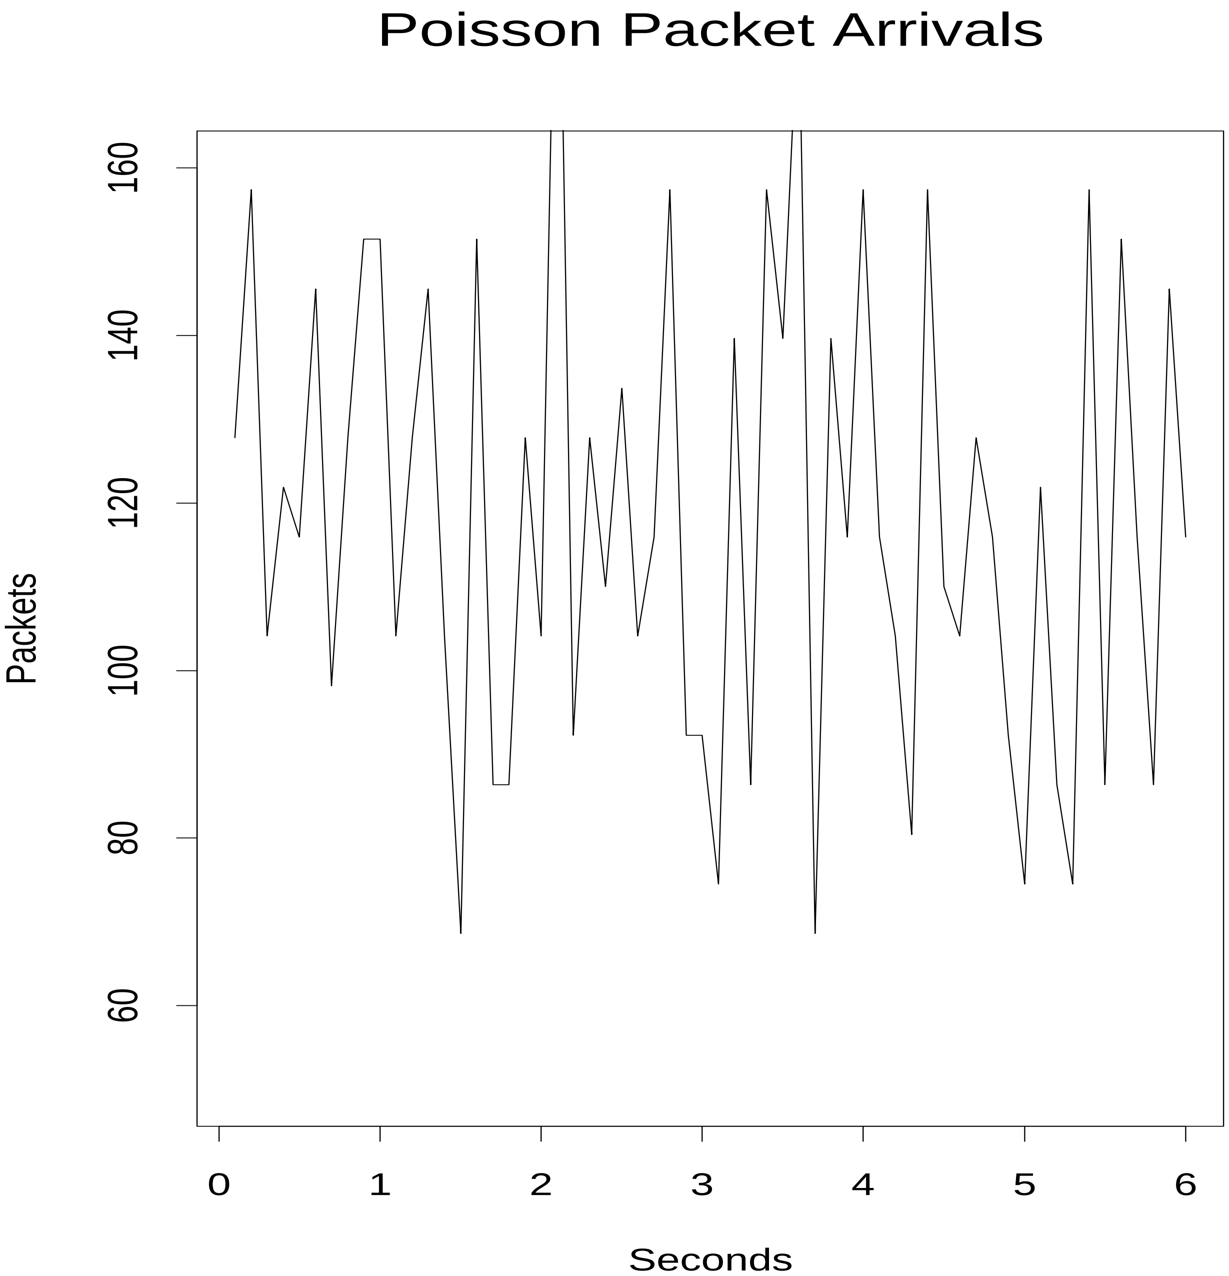
Poisson Packet Arrivals
Packets
60
80
100
120
140
160
0
1
2
3
4
5
6
Seconds

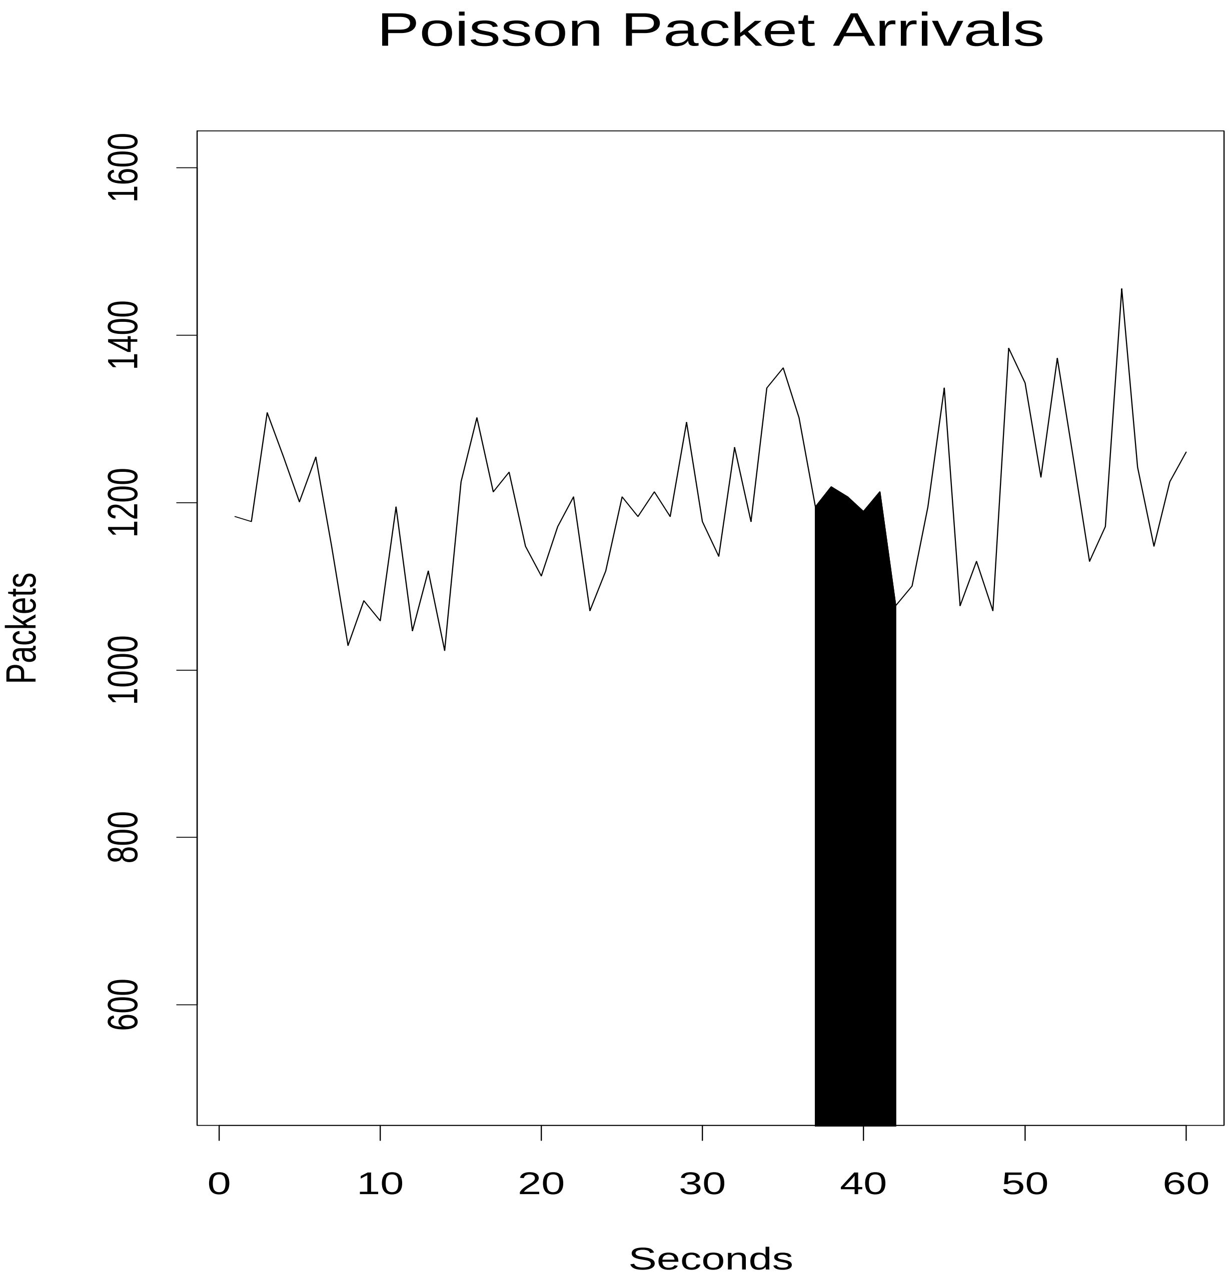
Poisson Packet Arrivals
Packets
600
800
1000
1200
1400
1600
0
10
20
30
40
50
60
Seconds

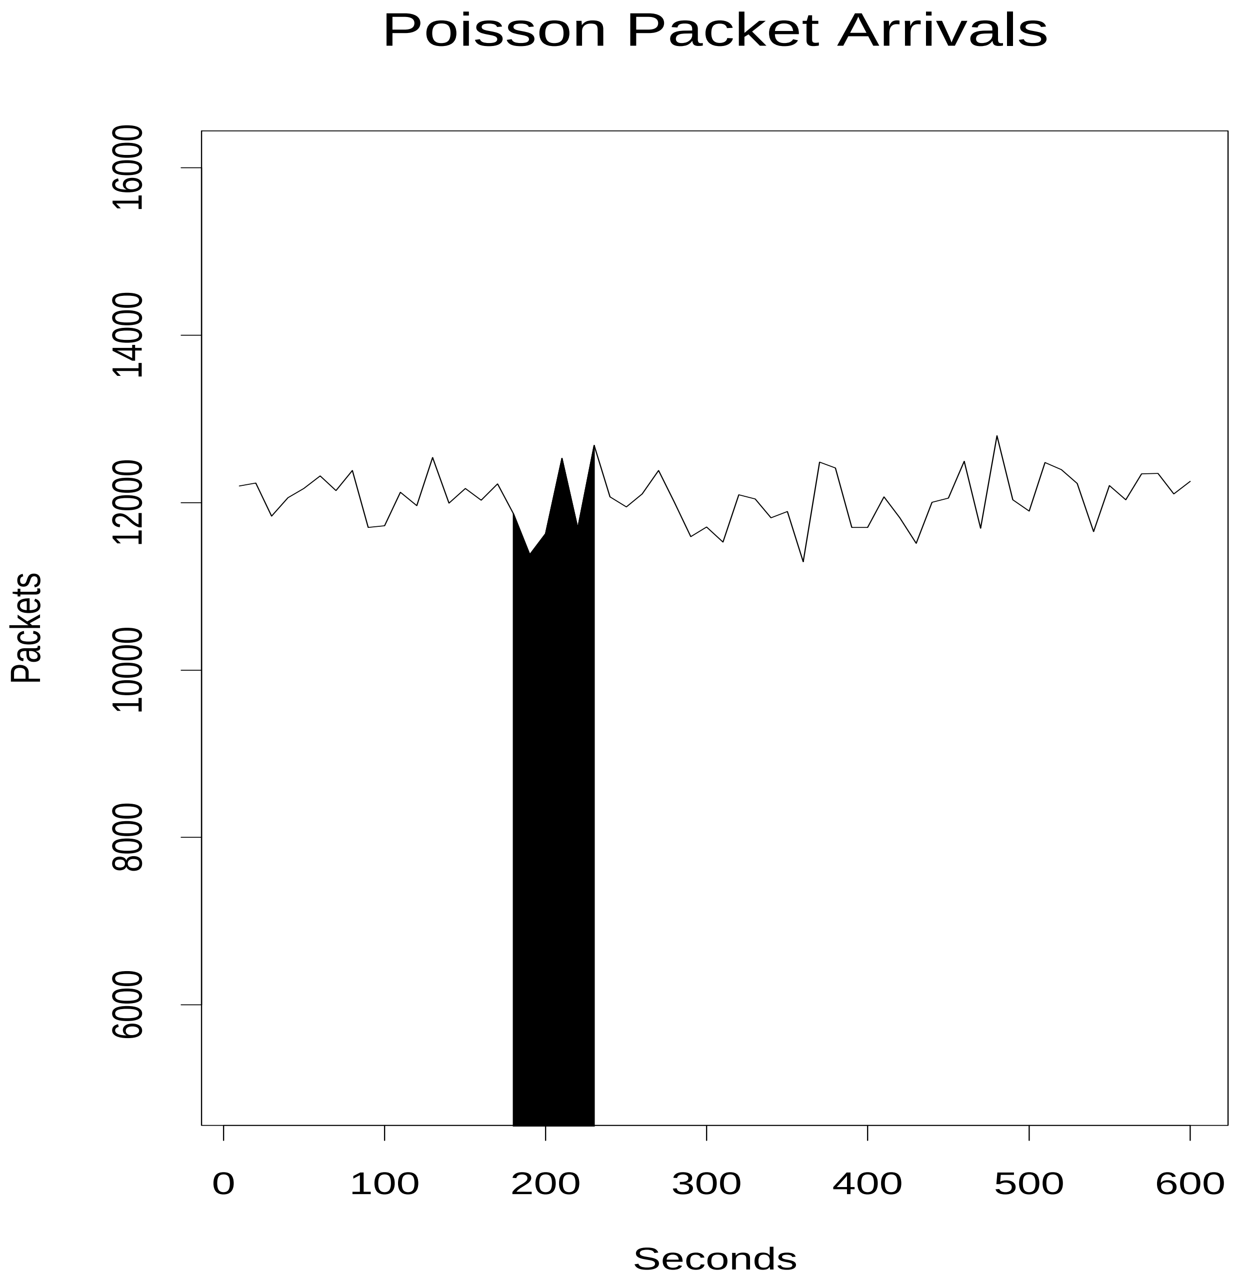
Poisson Packet Arrivals
16000
14000
12000
10000
8000
6000
Packets
0
100
200
300
400
500
600
Seconds

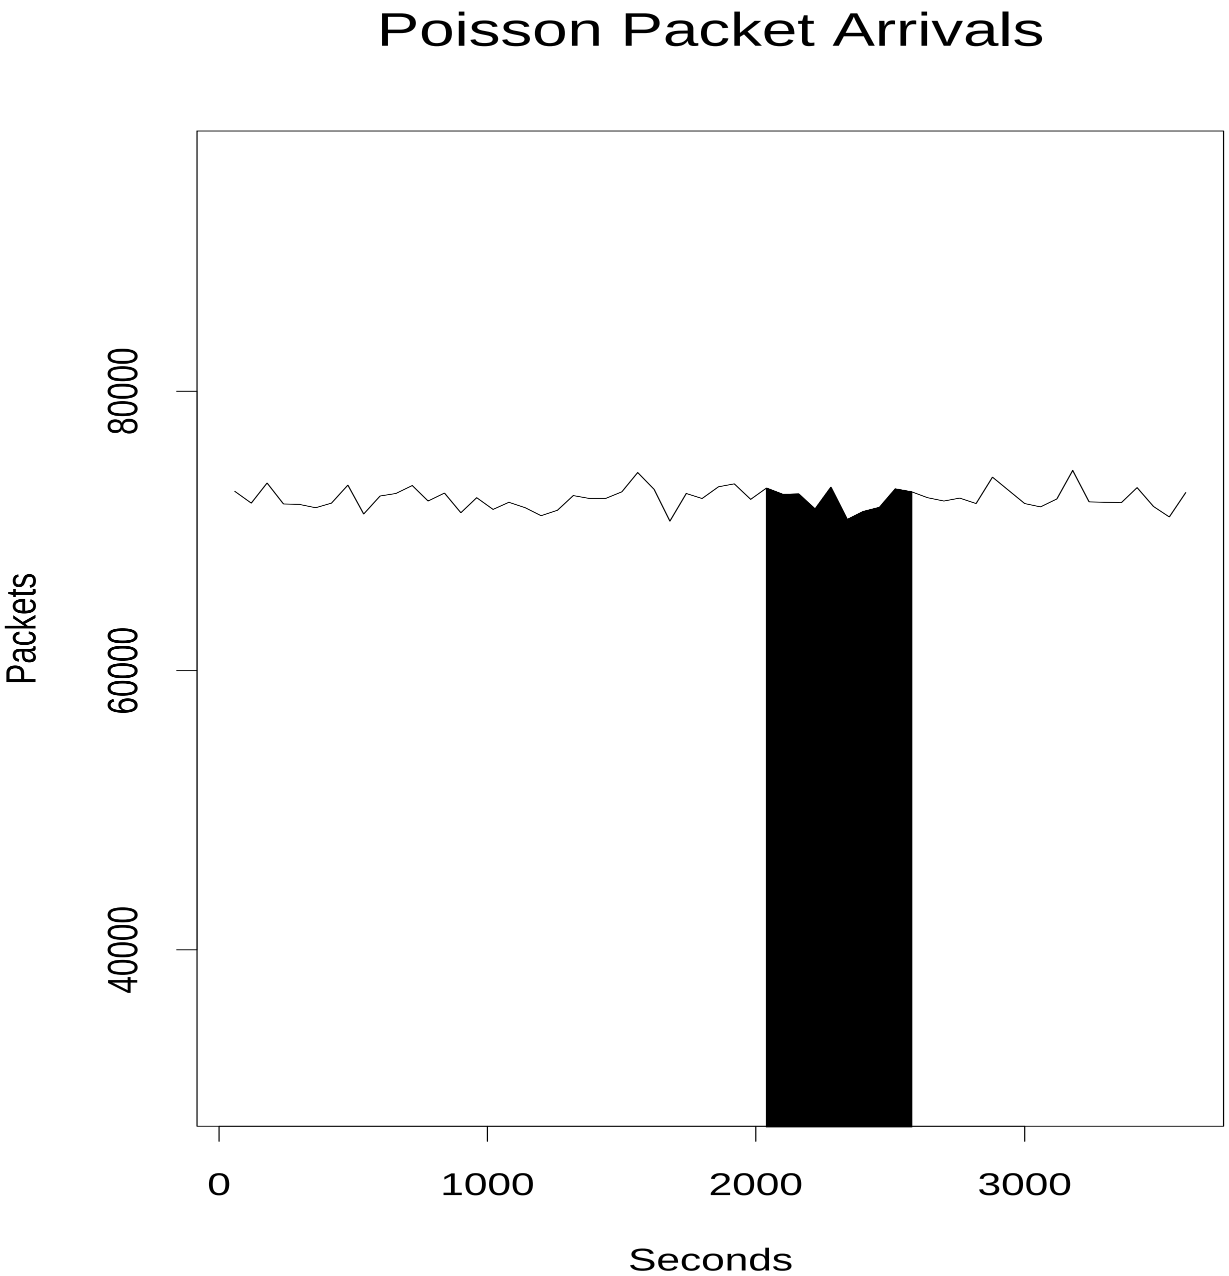
Poisson Packet Arrivals
Packets
40000
60000
80000
0
1000
2000
3000
Seconds

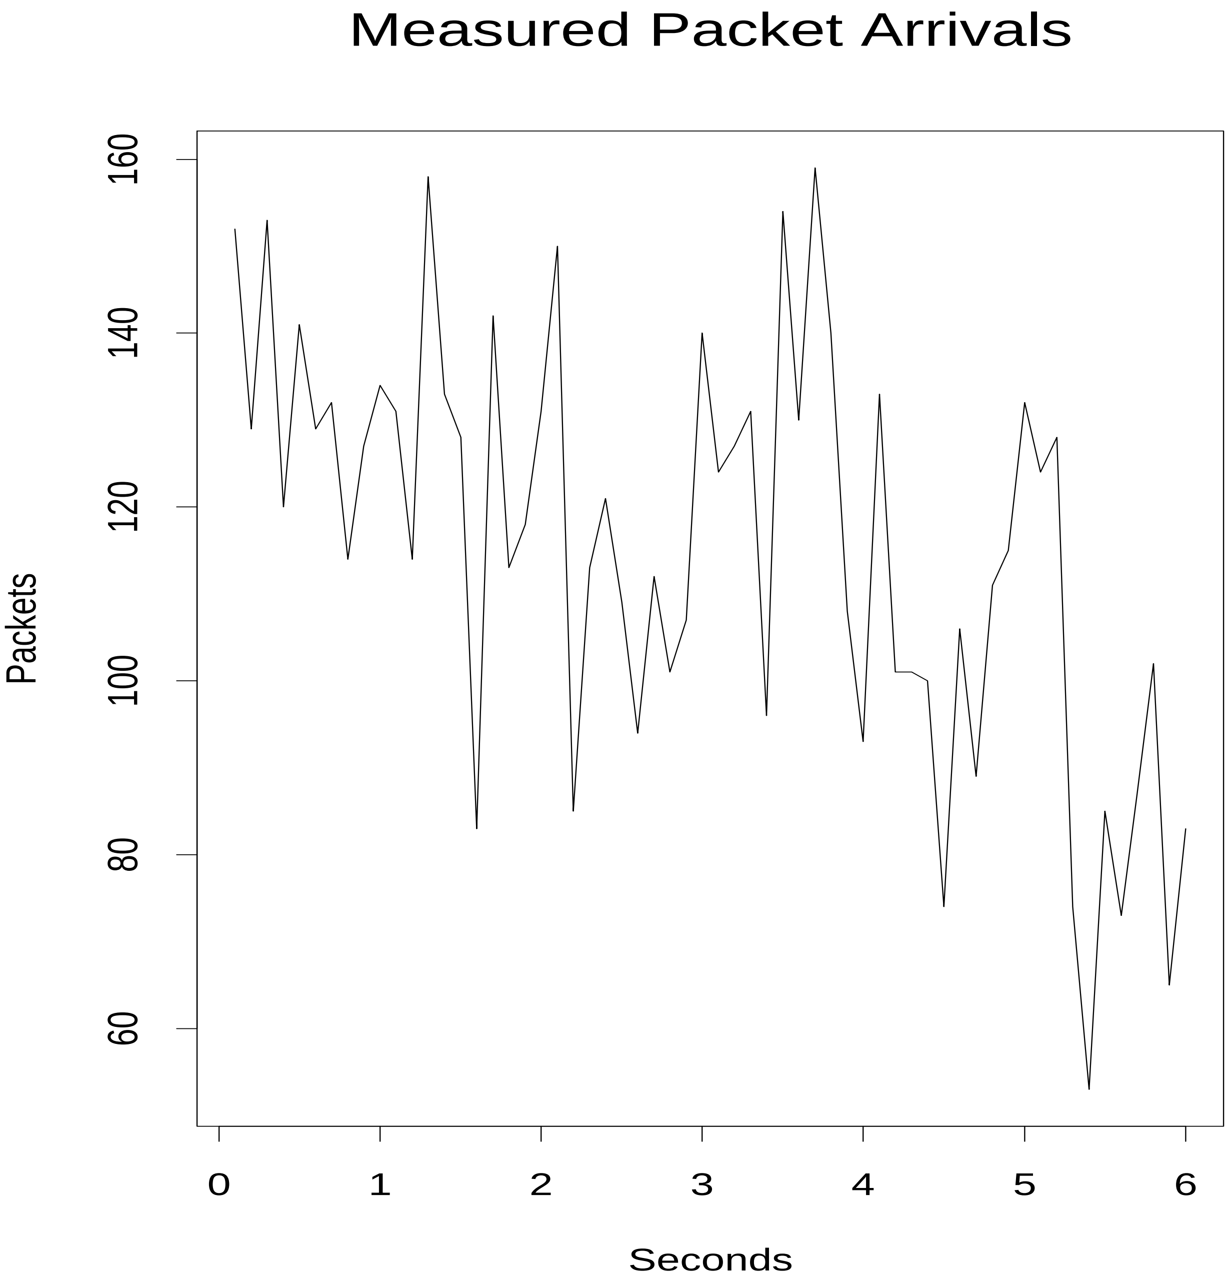
Measured Packet Arrivals
Packets
60
80
100
120
140
160
0
1
2
3
4
5
6
Seconds

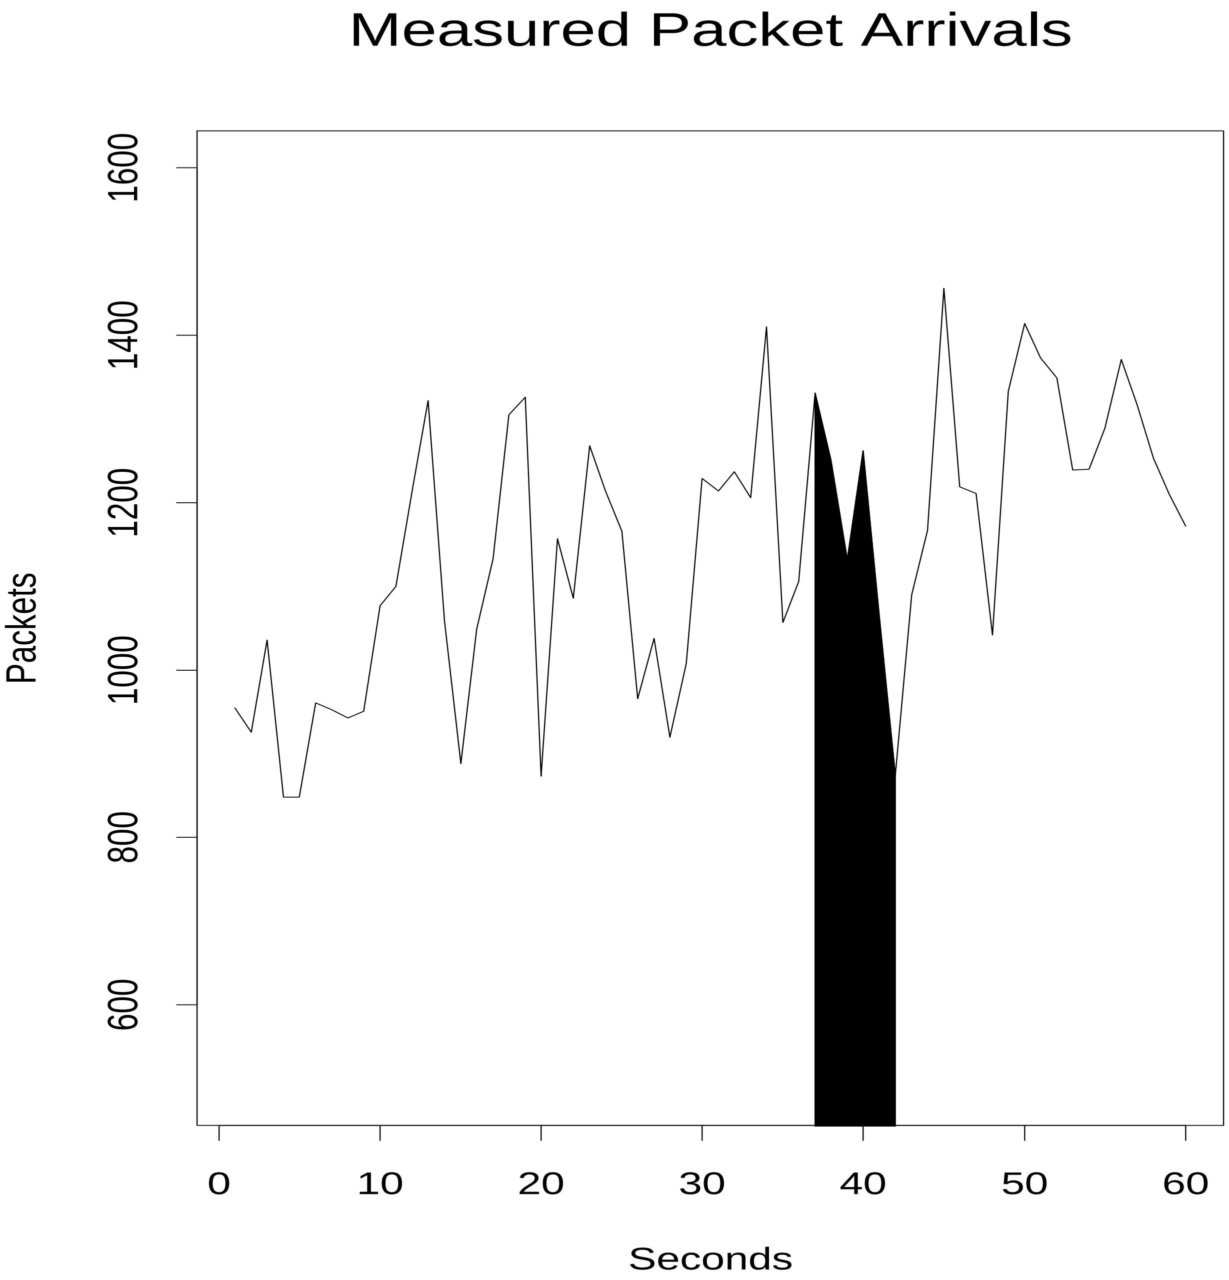

Measured Packet Arrivals
Packets
600
800
1000
1200
1400
1600
0
10
20
30
40
50
60
Seconds

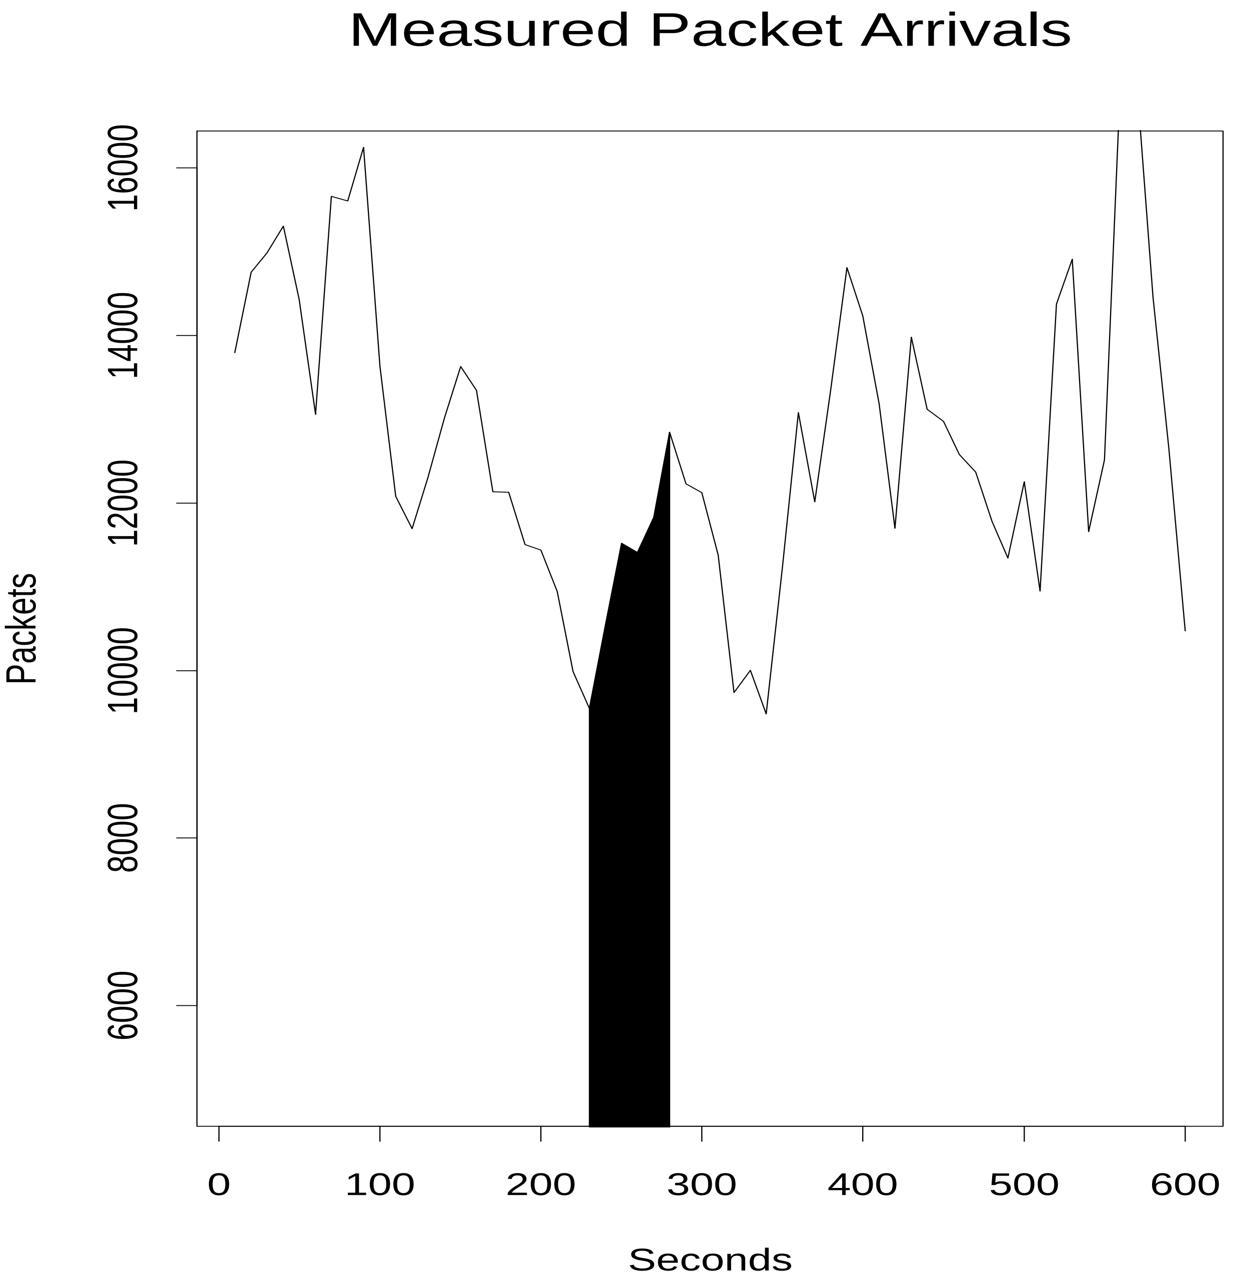
Measured Packet Arrivals
Packets
16000
14000
12000
10000
8000
6000
0
100
200
300
400
500
600
Seconds

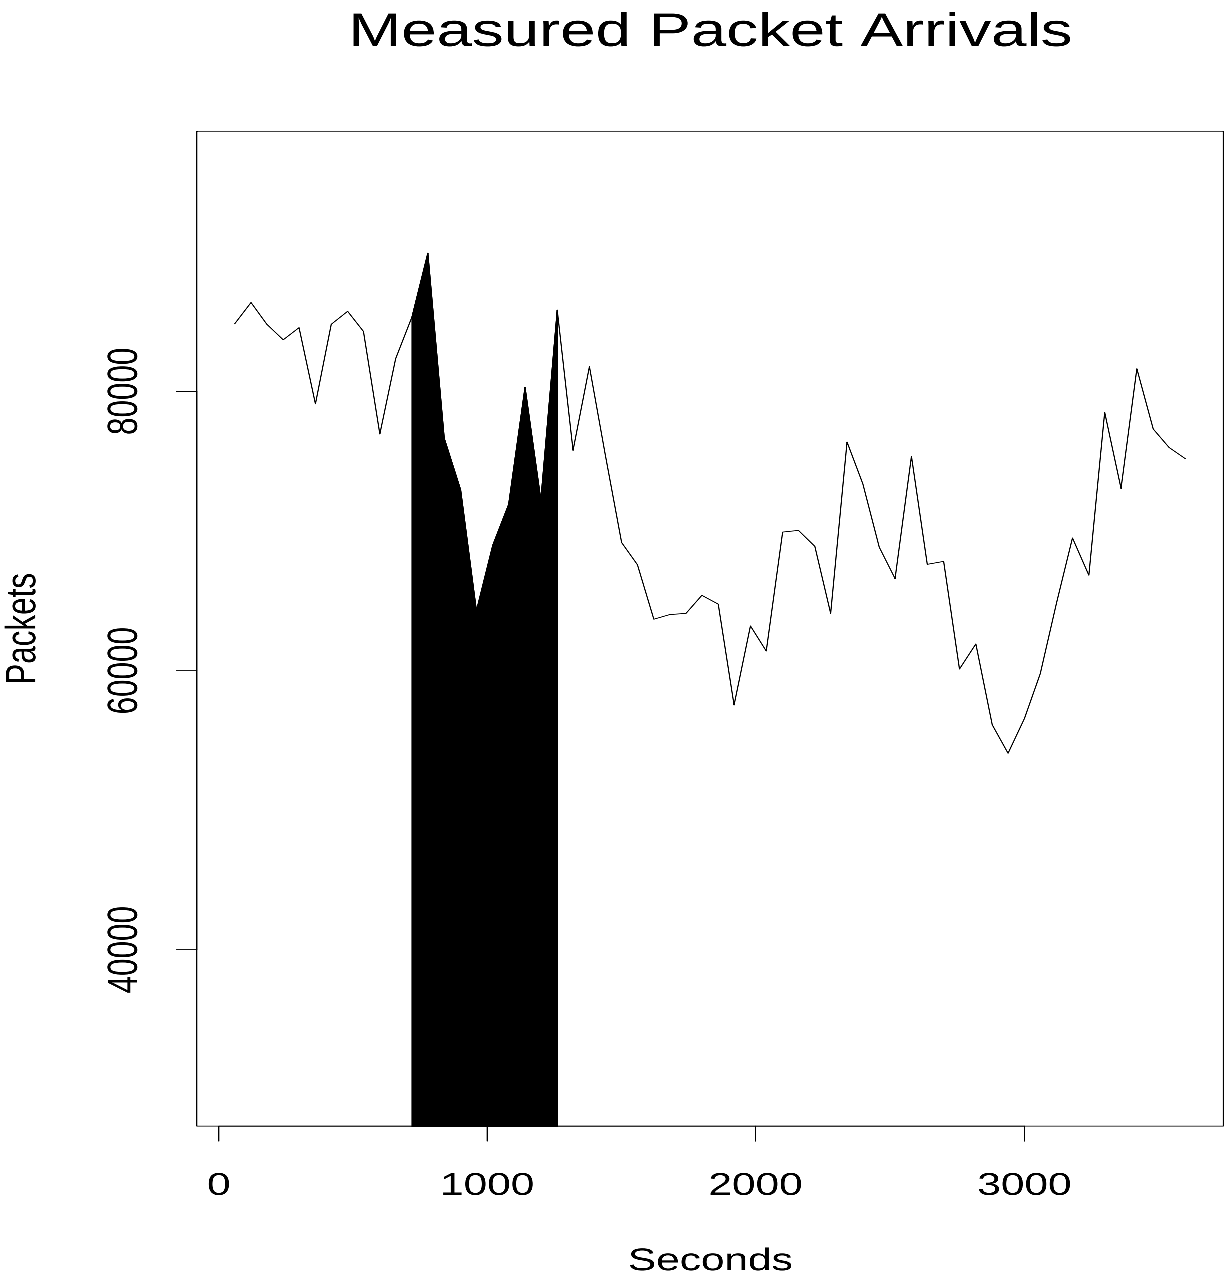
Measured Packet Arrivals
Packets
80000
60000
40000
0
1000
2000
3000
Seconds

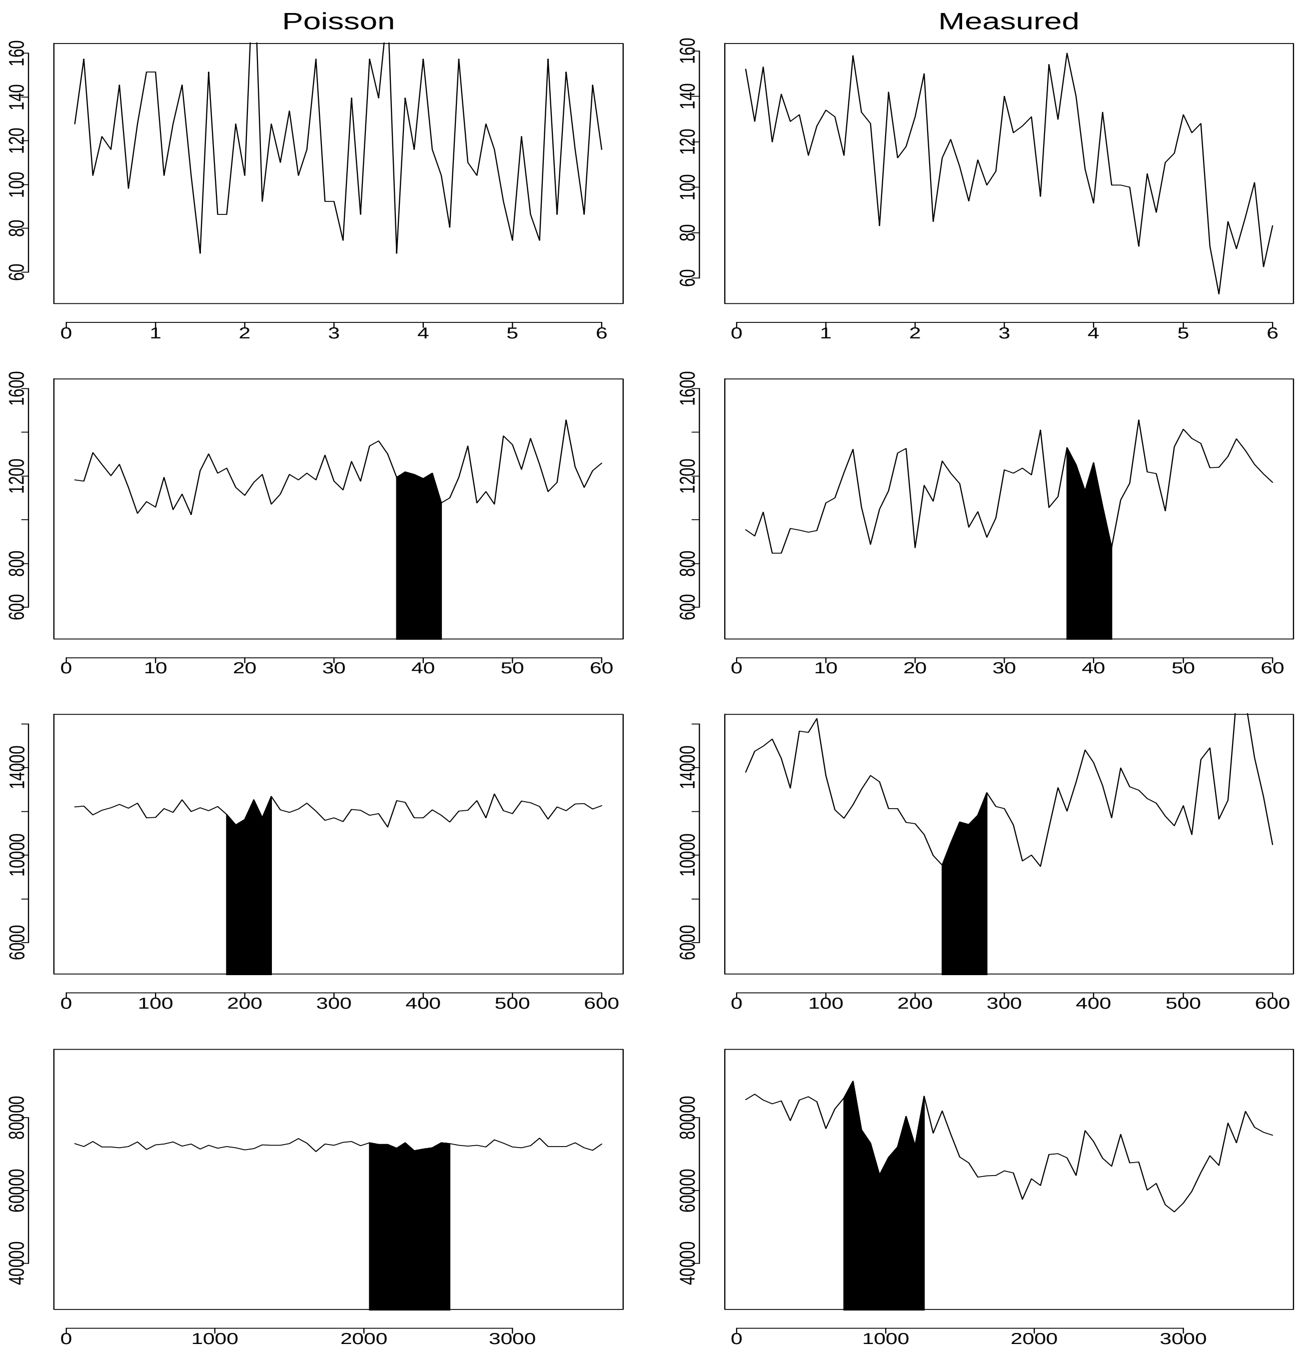

Poisson
Measured

Burstiness over many time scales:

Real traffic has strong, long-range correlations.

Power spectrum: flat for Poisson process, diverges to ∞ for $\omega \to 0$ in measurements.

$\Rightarrow$ To build Poisson-based models for this type of traffic takes major fiddling (many parameters).

$\Rightarrow$ But because the Internet is such a moving target, we are <u>desperate</u> for <u>parsimonious</u> models.

Measured DEC-WRL Traffic

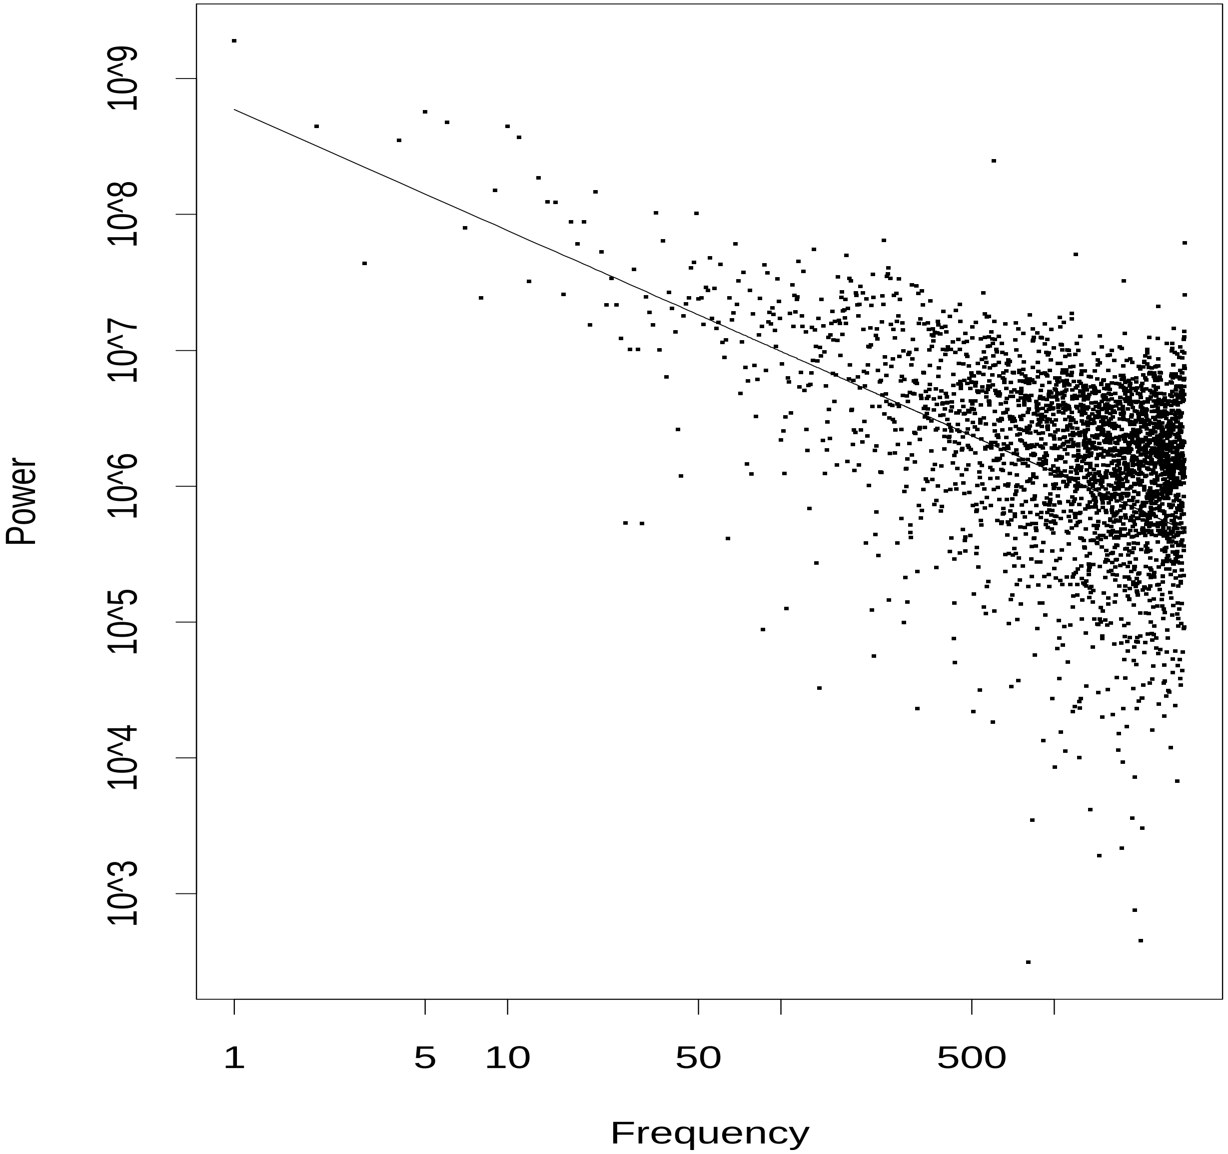

Fractal models:

In 1991, Fowler & Leland, looking at LAN traffic, observe that, in reality: *Traffic “spikes” (which cause actual losses) ride on longer-term “ripples”, that in turn ride on still longer-term “swells”*.

In 1993, Leland, Taqqu, Willinger & Wilson analyze additional LAN traffic. Propose fractal (self-similar; long-range dependent) models as parsimonious description of Ethernet traffic.

Bombshell! Within a year, fractal models emerged for Internet traffic, digitized video.

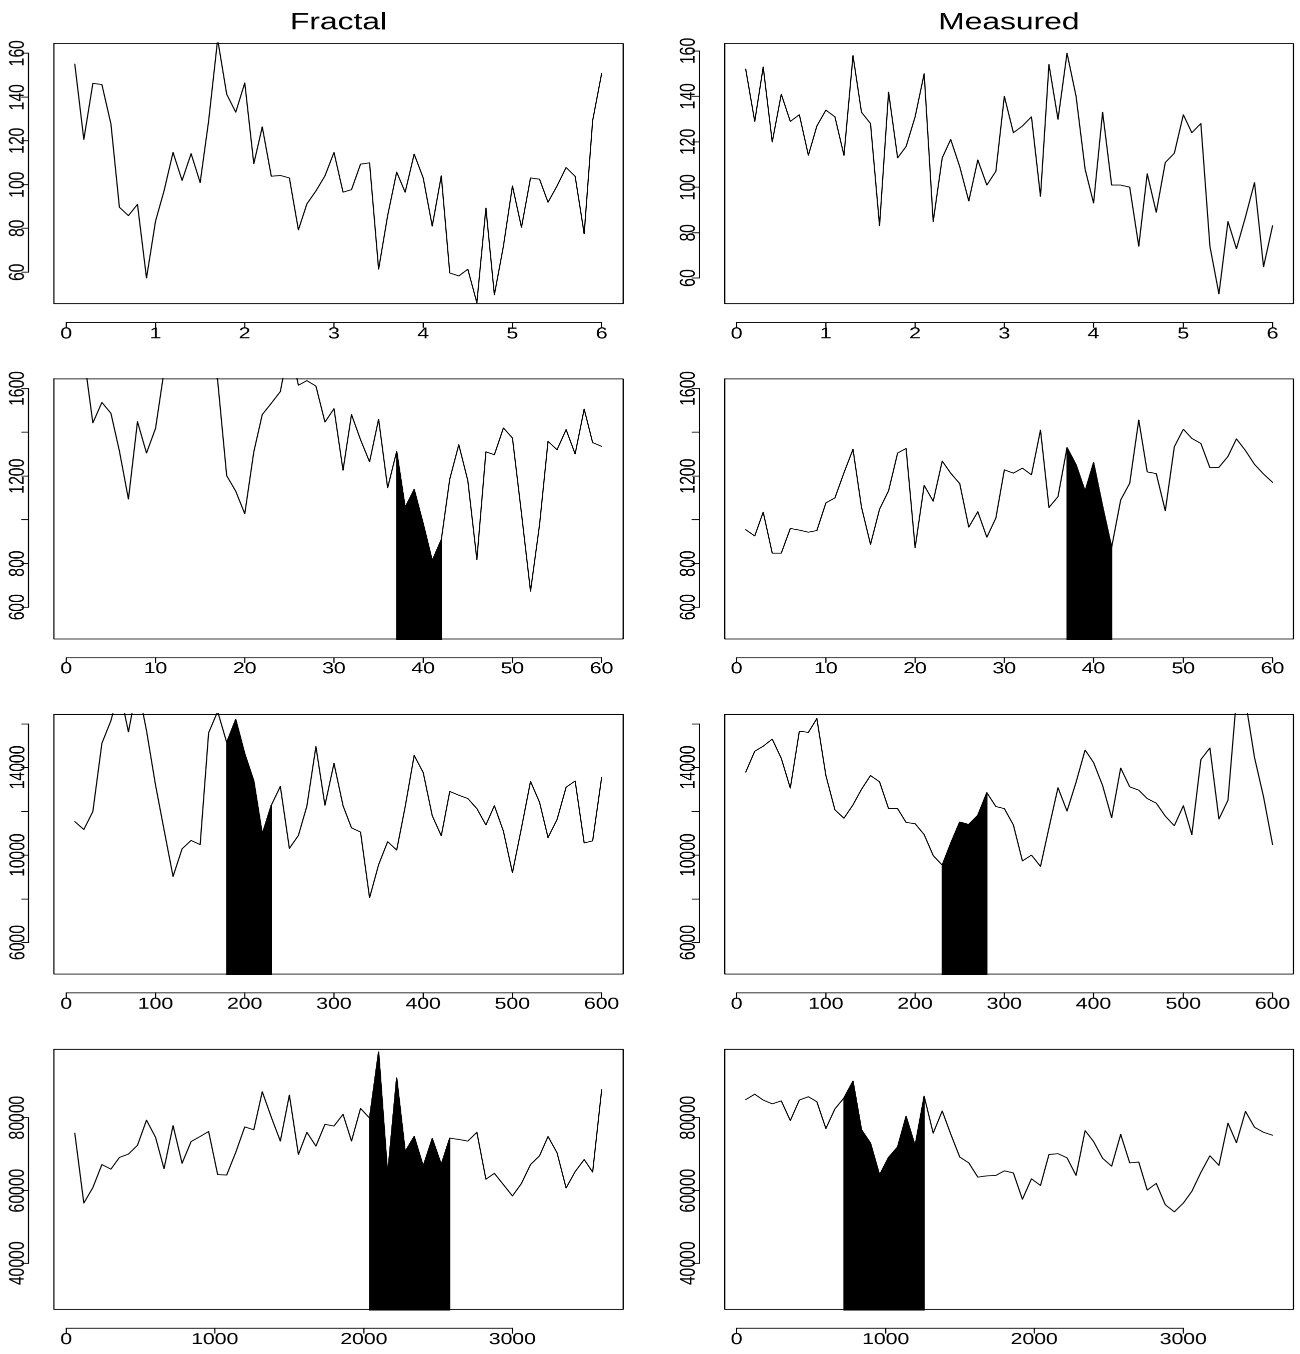
Fractal
Measured

Fractal models, con't:

Care taken to not be fooled by non-stationarity.

Applies on time scales of O(1 sec) to O(1 hour).

Maybe not <u>exact</u>, but certainly a <u>good approximation</u>.

Major implication: queue delays / drop probabilities *much* higher than otherwise expected.

Fractal models, con't:

Evidence that finer time scales well-modeled using *multi-fractals*.

Fledgling explanations how this relates to *congestion control*.

Wavelet analysis the tool of choice.

Looking at aggregate traffic gives insight for individual traffic!

Heavy tails:

1993 prediction: fractal traffic behavior can be explained if distribution of *individual* activity periods has *infinite variance*.

Pareto distribution: $\overline{F}(x) = (x/a)^{-\beta}$.
If $\beta < 2$, then *infinite variance*.

Can test for Pareto distribution by plotting $\log \overline{\widehat{F}}_n(x)$ vs. $\log x$.

Straight line $\Rightarrow$ Pareto. Slope estimates $-\beta$.

Web connection durations (226,386 observations):

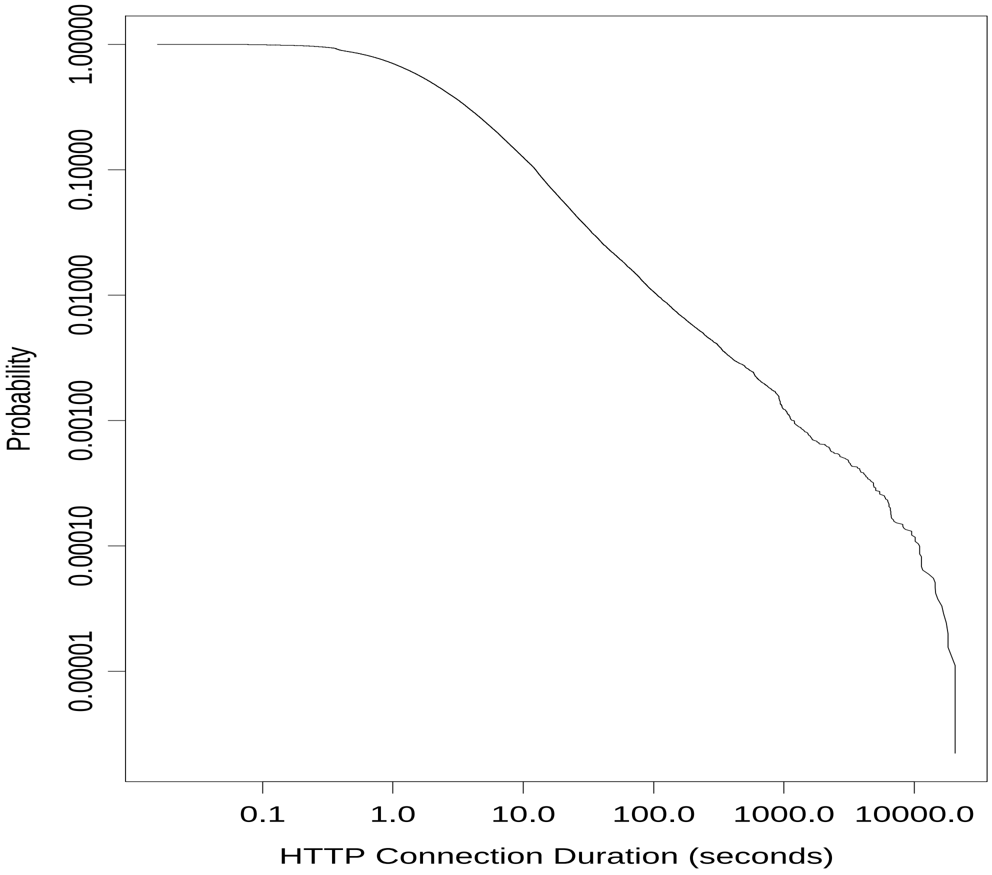

Web connection duration tail (28,469 observations):

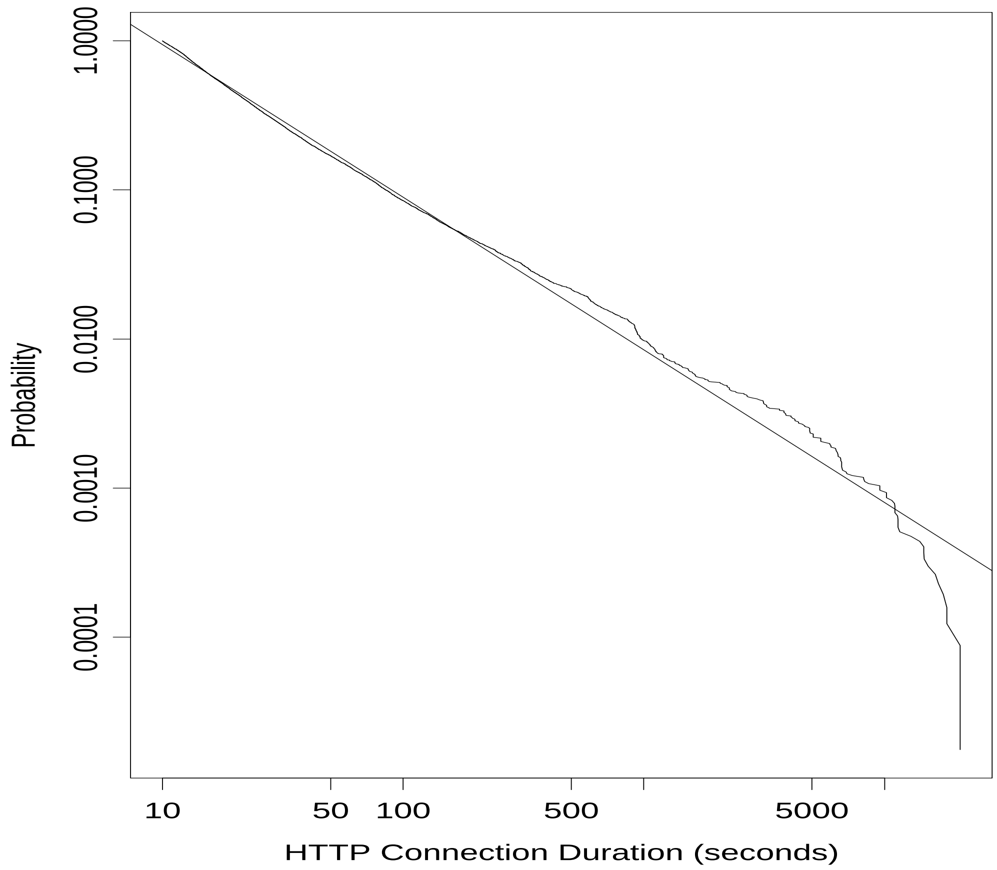

$\beta \approx 1.0$. Infinite variance.

Infinite mean?

Web connection durations:

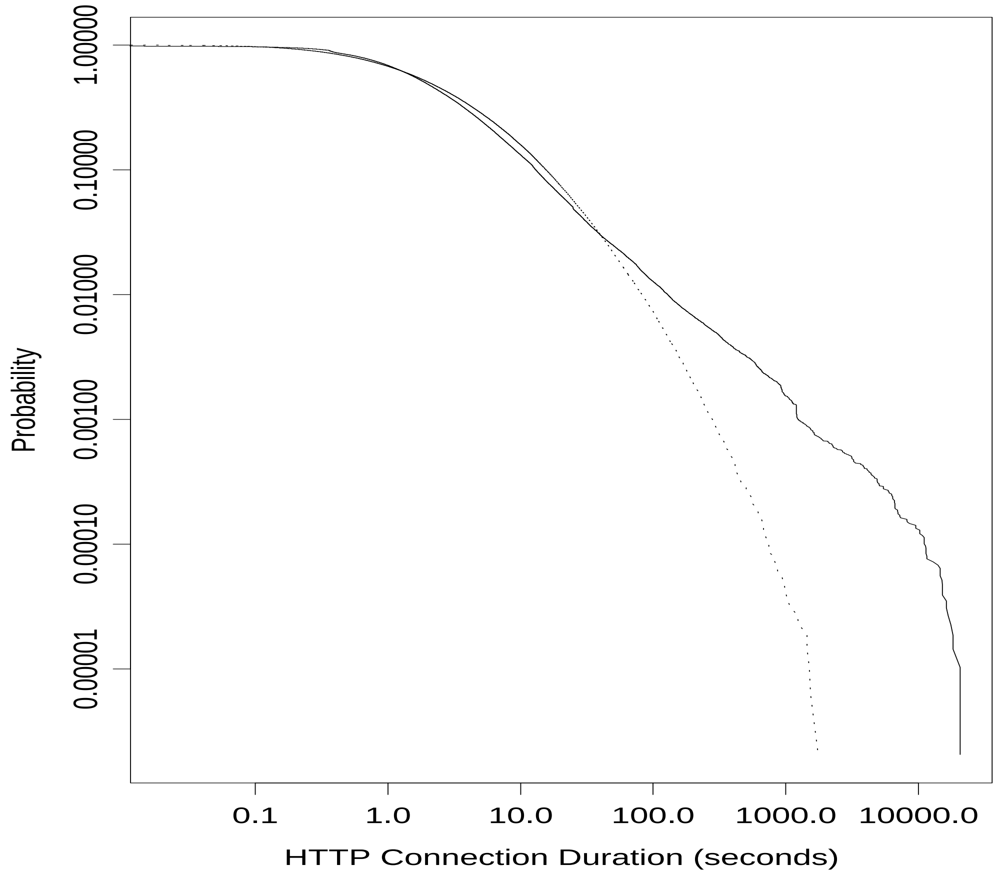

Major difference in upper tail vs. fitted log-normal distribution.

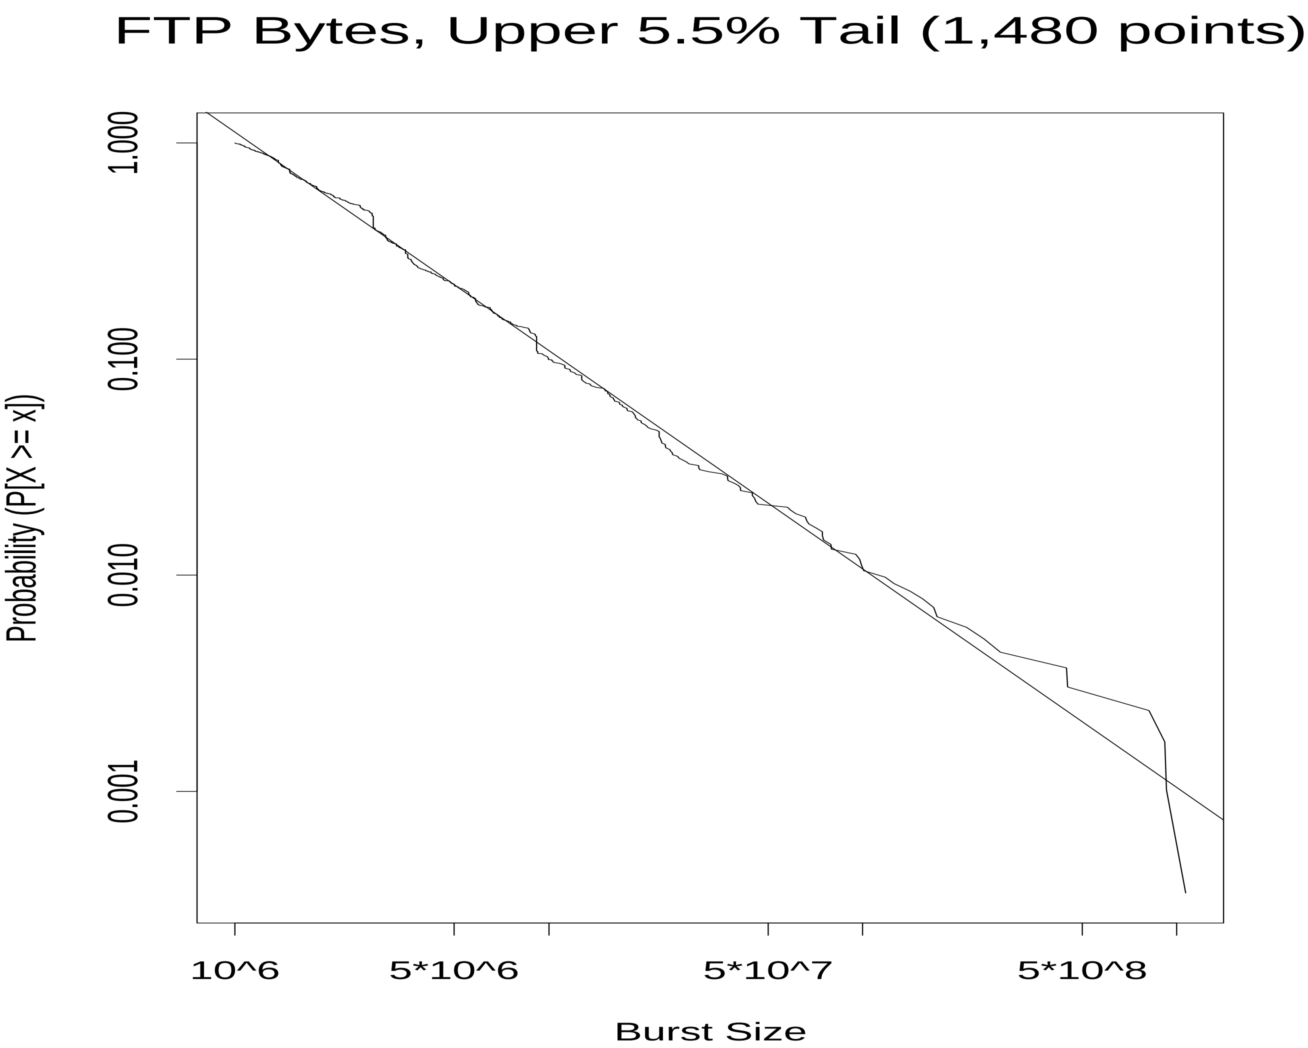
FTP Bytes, Upper 5.5% Tail (1,480 points)
Probability (P[X >= x])
1.000
0.100
0.010
0.001
10^6
5*10^6
5*10^7
5*10^8
Burst Size

Heavy tails in computing:

CPU time consumed by Unix processes [LO86,H-BD96].

Size of Unix files [Ir93].

Size of compressed video frames [Ga93].

Size of FTP bursts [PF94].

Telnet packet interarrivals [D+91,PF94].

Size of Web items [CB96].

<u>Ethernet bursts</u> [W+95].

Summary:

Immense diversity:

- must analyze across multiple datasets
- search for “laws” (invariants), not exact models

One such invariant: heavy tails in size distributions.

Another: strong, pervasive long-term correlations:

- accurate estimation tricky
- we're serious about fractal modeling

$\Rightarrow$ Only way we know how to build viable, large-scale Internet systems is to make them very very robust.