An Overview of
Internet Measurement and Modeling

Vern Paxson

Network Research Group
Lawrence Berkeley National Laboratory
EECS Division, U.C. Berkeley
vern@ee.lbl.gov

March 20, 1996

Overview:

Why Internet measurement is hard.

Empirical and analytic models of TCP connections.

The failure of Poisson modeling.

The advent of fractal modeling.

Internet questions:

1. How many people use the Internet?
2. What does a typical connection look like?
3. What's the dominant application?
4. How often are data packets dropped?
5. How big are the queues inside the network?
6. How can I model (or simulate) it faithfully?

How can I answer questions like these?

Internet answers:

1. Basically, no one has a clue how many users there are. The number of hosts, however, is periodically measured.

2. Connections are so diverse “typical” is meaningless.

3. The dominant application depends on what you're concerned about. Right now: HTTP. This will for sure change.

4, 5. No one knows how often packets dropped or queued.

6. Faithful Internet modeling is an open research problem.

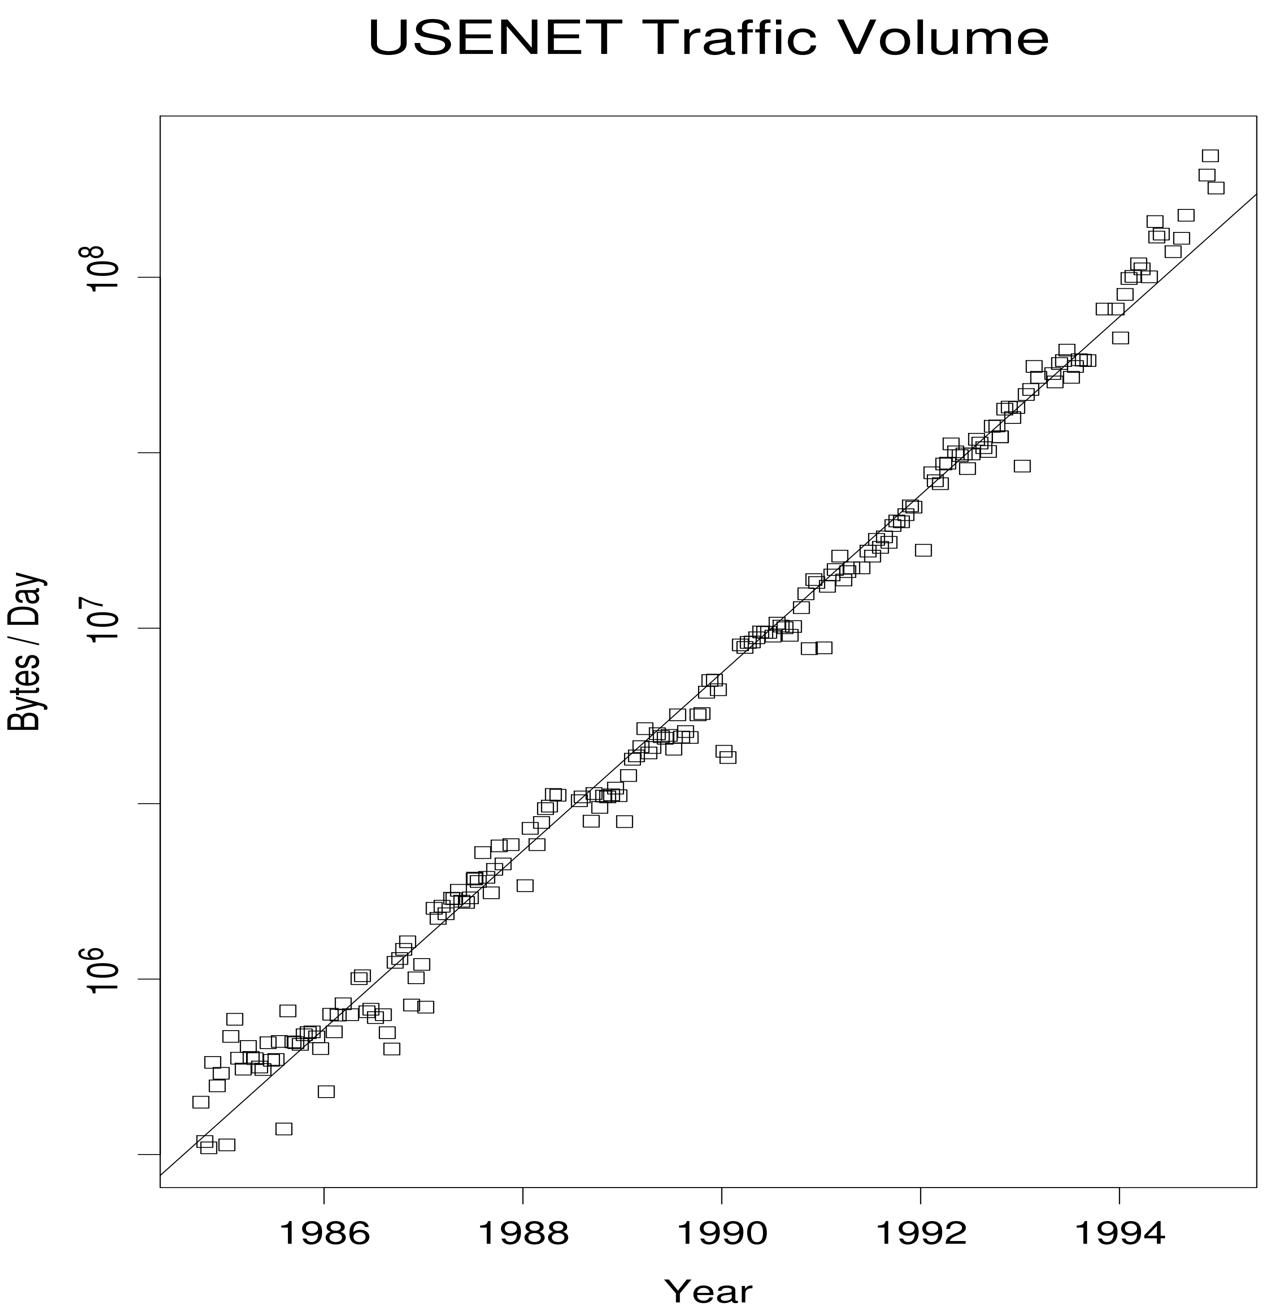
USENET Traffic Volume
Bytes / Day
10^6
10^7
10^8
1986
1988
1990
1992
1994
Year

The Internet — huge & diverse:

In January, 1990, the Internet had about 200,000 hosts. Six years later, it had about 9,500,000.
$\Rightarrow$ If you blink, you're out of date.

In January, 1991, about 15% of the data bytes going into LBNL were Network News. At USC, it was 60%.
$\Rightarrow$ There's no such thing as a typical site.

In October, 1992, the median FTP data transfer at LBNL was 4,500 bytes. Five months later, it was 2,100 bytes.
$\Rightarrow$ There's no such thing as a typical time period.

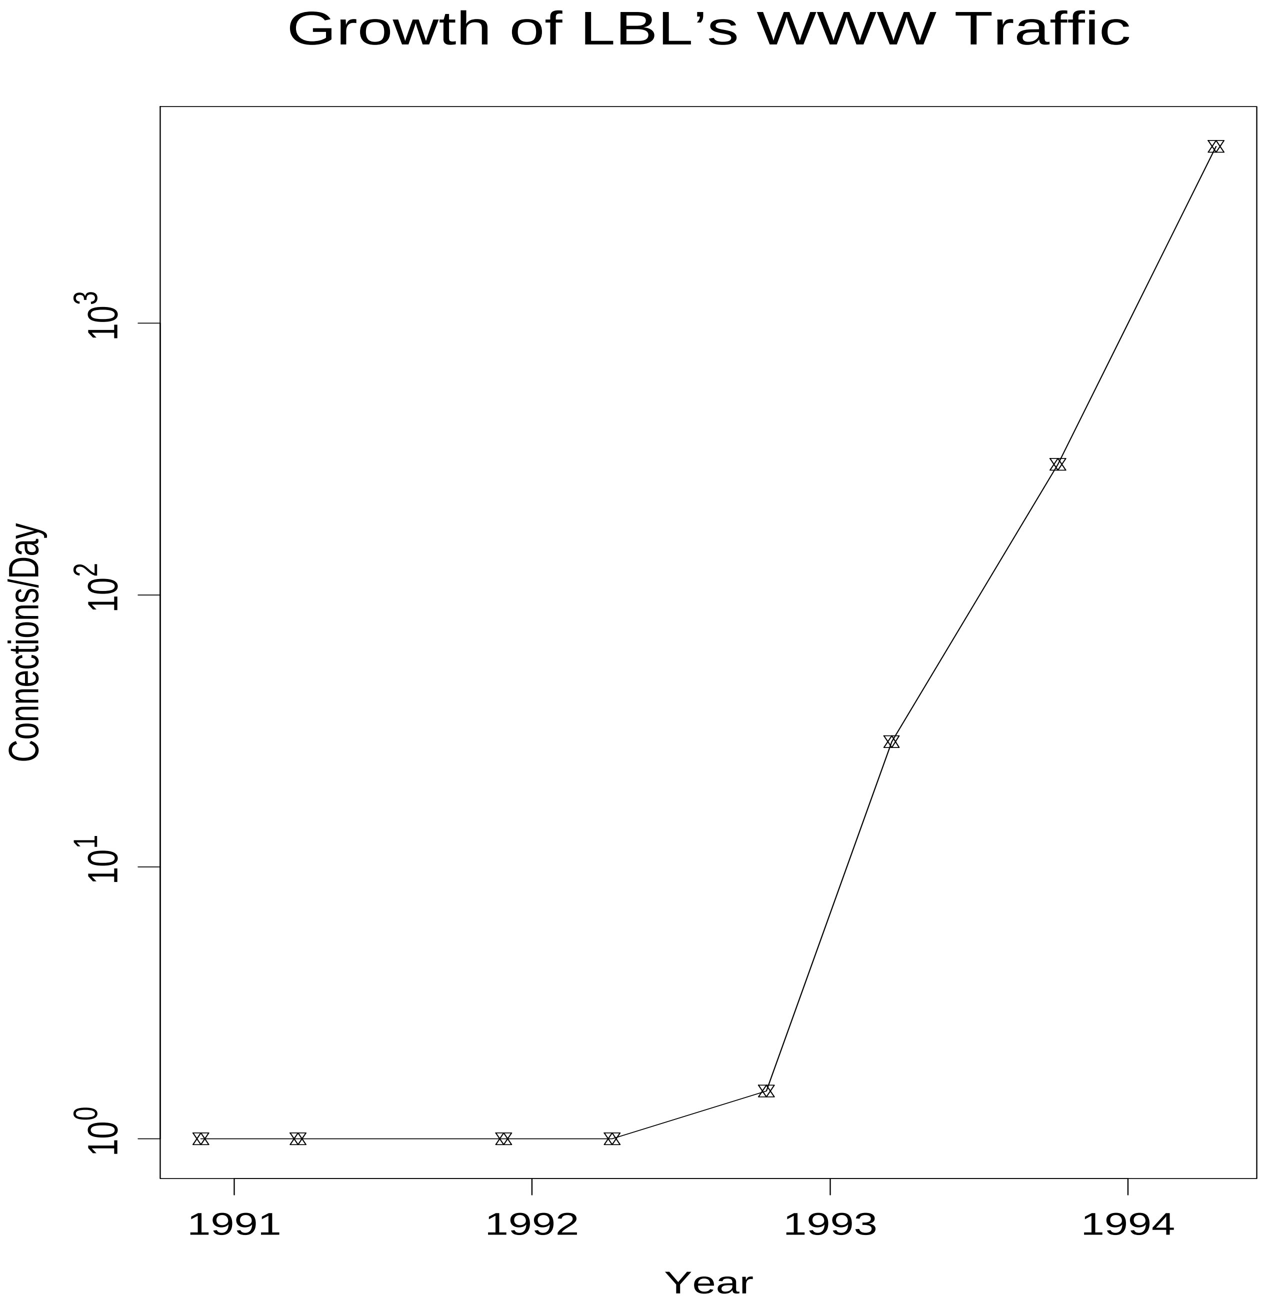

Growth of LBL's WWW Traffic
Connections/Day
10^0
10^1
10^2
10^3
1991
1992
1993
1994
Year

New applications radically alter the landscape:

In the *month* of October, 1992, LBNL made *45* World-Wide Web connections. Three and a half years later, it makes *200,000* 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 LBL, and more than 50% at Bellcore.

Traditional network modeling — Poisson traffic

Idea: everything is independent of everything else. There's just a single **rate** of calls/hour.

$\Rightarrow$ **memoryless** property.

Times between calls ("interarrivals") distributed exponentially: $F(x) = 1 - e^{-\lambda x}$, and *independently*. Exponential is the *memoryless* distribution.

Number of calls ("arrivals") has Poisson distribution:

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

Memoryless property makes analysis easy:

Nothing depends on what happened previously! Or, throw in a queue, and you introduce “short range dependence” — things influence the near-term. $\Rightarrow$ *Analysis is easy.*

There are many variants of Poisson modeling, but all share same basic properties.

Poisson modeling a huge success: telephone system works because of it. $\gg 1,000$ papers published in last 30 years.

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.”

Traditional/Poisson assumptions:

Independence among arrivals.

(for sessions, connections, packets)

Exponentially distributed sizes, durations.

Bulk transfers are infinitely large, and unidirectional.

Interactive connections transfer equal amounts of traffic in each direction.

Virtually all of these are wrong.

It took 25 years to figure this out—a testament to the power of analytic principles in shaping our thinking.

The USC/UCB/Bellcore study:

What they measured: all packet headers on the DMZ networks connecting USC, UCB, and Bellcore to the Internet.

When: 1991 / 1989 / 1989.

Duration: 1 / 1 / 3 days.

Drops: 0.6% / 0 / 0.

Clock resolution: 20 msec / 10 μ sec / 10 μ sec.

Volume: 5.2 Mpkt / 5.9 Mpkt / 1.7 Mpkt.

Identifying connections: delimited by 20-minute silences.

The USC/UCB/Bellcore study - what they found:

Great variation in traffic mix between the sites.

Up to 90% of bulk transfers send $<$ 10 KB, request/response.

50% of interactive connections last $<$ 90 s. The responder can generate 10 times more data than the originator.

Telnet packet interarrivals have a fixed distribution.

Most bulk transfers are bidirectional.

The USC/UCB/Bellcore study - Tcplib:

Data used to construct *empirical* models.

Step one: for a random variable (e.g., Telnet packet interarrival T), construct $\widehat{F}(t)$ from dataset: $\widehat{F}(t) = P(T \leq t)$.

Step two: to simulate interarrivals, generate uniform random numbers and use *inverse-transform method* to distribute using empirical distribution: $T = \widehat{F}^{-1}(U)$ for $U \sim$ Unif(0,1).

No models for bulk transfer interarrivals or rates: *Why?*

Empirically-derived analytic models:

What I measured: seven months of LBL wide-area SYN/FIN data, 1990-1993 (2.5 million connections).

What others donated: a single day from UCB, USC, coNCert, UK–US; 3 days from DEC-WRL; 13 days from Bellcore (350 thousand connections). 1991, 1989.

Identifying connections: SYN establishment through FIN termination.

How to generalize 1 GB of trace data into useful models?

Empirical vs. analytic models:

	Empirical	Analytic
Basic model	raw, measured distribution	analytic function fit to distribution.
Form	$O(1000)$ numbers	$F(x;\theta)$
Example	Tcplib (Danzig et.al.)	Log-normal
Pros	Exactly represents an aspect of the Real World. Preserves quirks (if present).	Mathematically tractable. Concise. Easy to communicate. Easy to compare.
Cons	No good if atypical.	No good if oversimplifies.

Is a model any good?

Goodness of fit tests: look for “statistically exact” fits. [DaS86]

No such luck.

And not surprising.

(26,306 throws of 12 dice . . . [M74])

Measure “discrepancy”: how closely does one distribution match another? λ^2 measure [PJ90] (derived from χ^2). Values can be meaningfully compared.

Test fit of analytic models vs. empirical.

Test fit of single model to different datasets.

Bytes generated by Telnet responder:

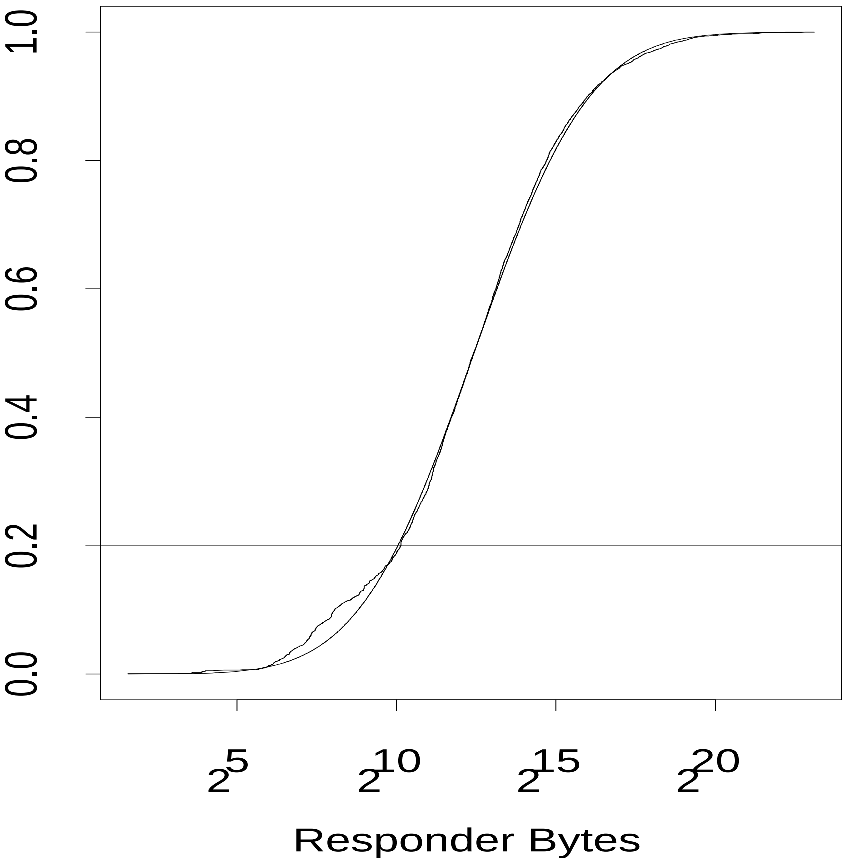

Goodness of fit test: failed!

Summary of models:

Proto.	Variable	Model	Corr.
telnet	originator bytes (T_{orig})	log-extreme	$+$
	responder bytes (T_{resp})	lgnorm, 80-100%	$+$
	duration secs. (T_{dur})	lgnorm	$+$
nntp	originator bytes (N_{orig})	lgnorm	$++$
smtp	originator bytes (S_{orig})	$2\times$ lgnorm + 300B	$++$
ftp	connection bytes (F_{conn})	lgnorm	$++$
	session bytes (F_{sess})	lgnorm	no
	burst bytes (F_{burst})	Pareto, 95-100%	$++, \pm$

Distribution of nntp originator bytes:

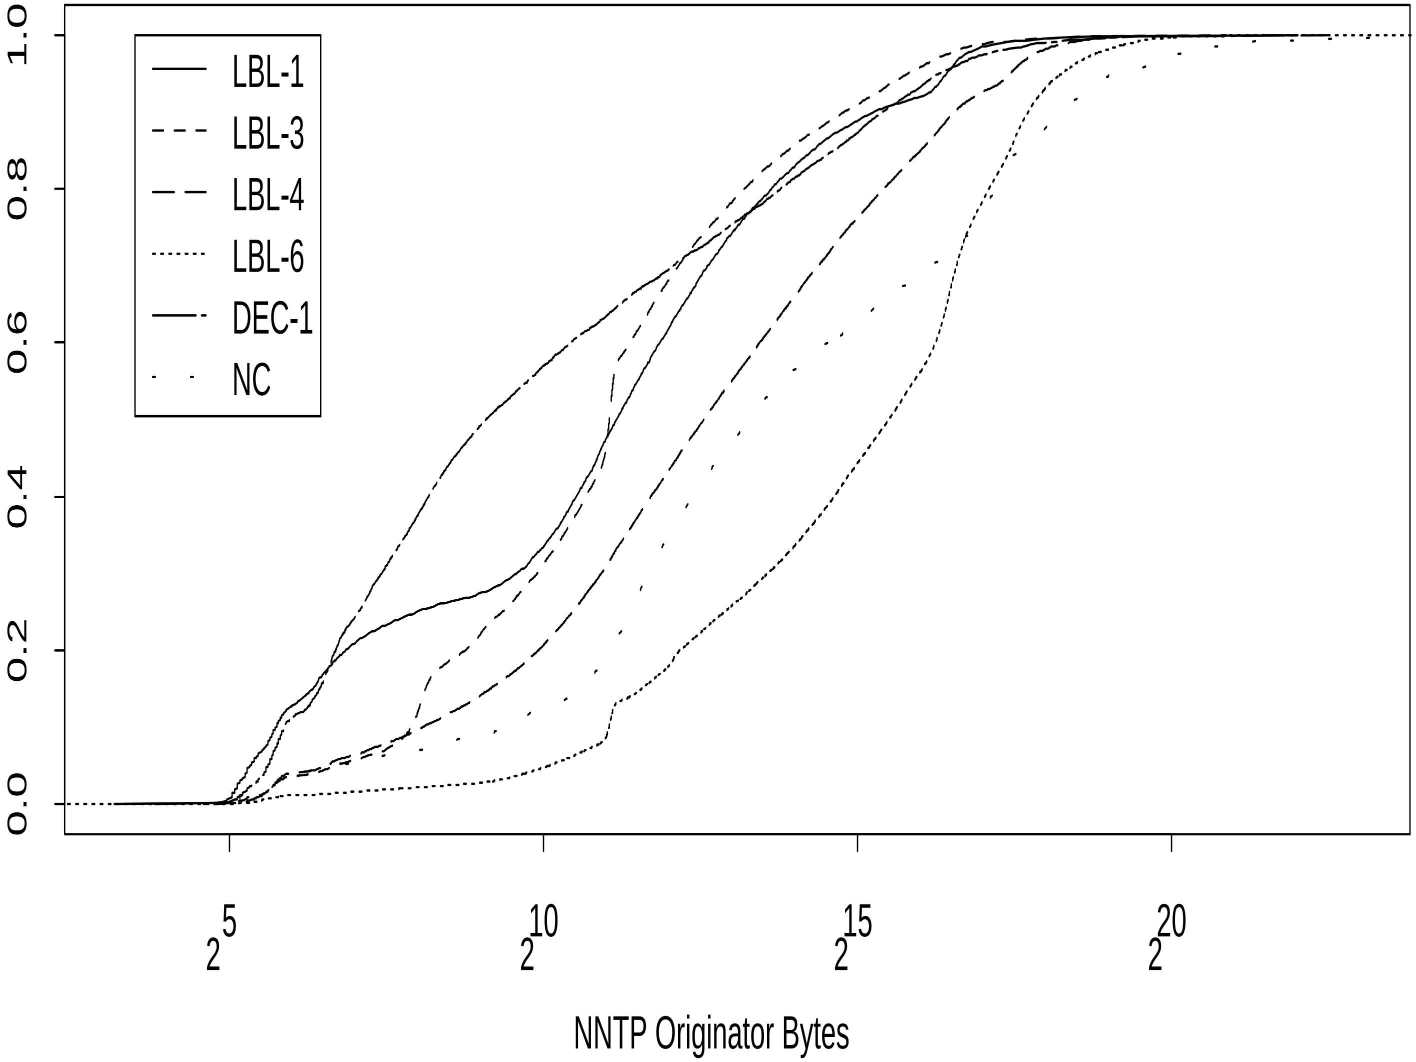

Findings:

1. Analytic models describe traffic as well as empirical ones.
2. Both types of models poorly describe the upper 1% tails.
3. Traffic varies significantly, both over time site-to-site.
4. Use log-normal distributions for size of bulk-transfer traffic.
5. Bulk-transfer traffic is not strongly bidirectional.

Findings, con't:

6. The computer side of a TELNET connection sends 20 times as many bytes as the user side.

7. Half of all* FTP *sessions* transfer more than 32 KB; 1/6 transfer more than 500 KB.

8. The upper tail of FTP "bursts" is so large that 2% of the bursts account for 50-80% of <u>all</u> the FTP data bytes.

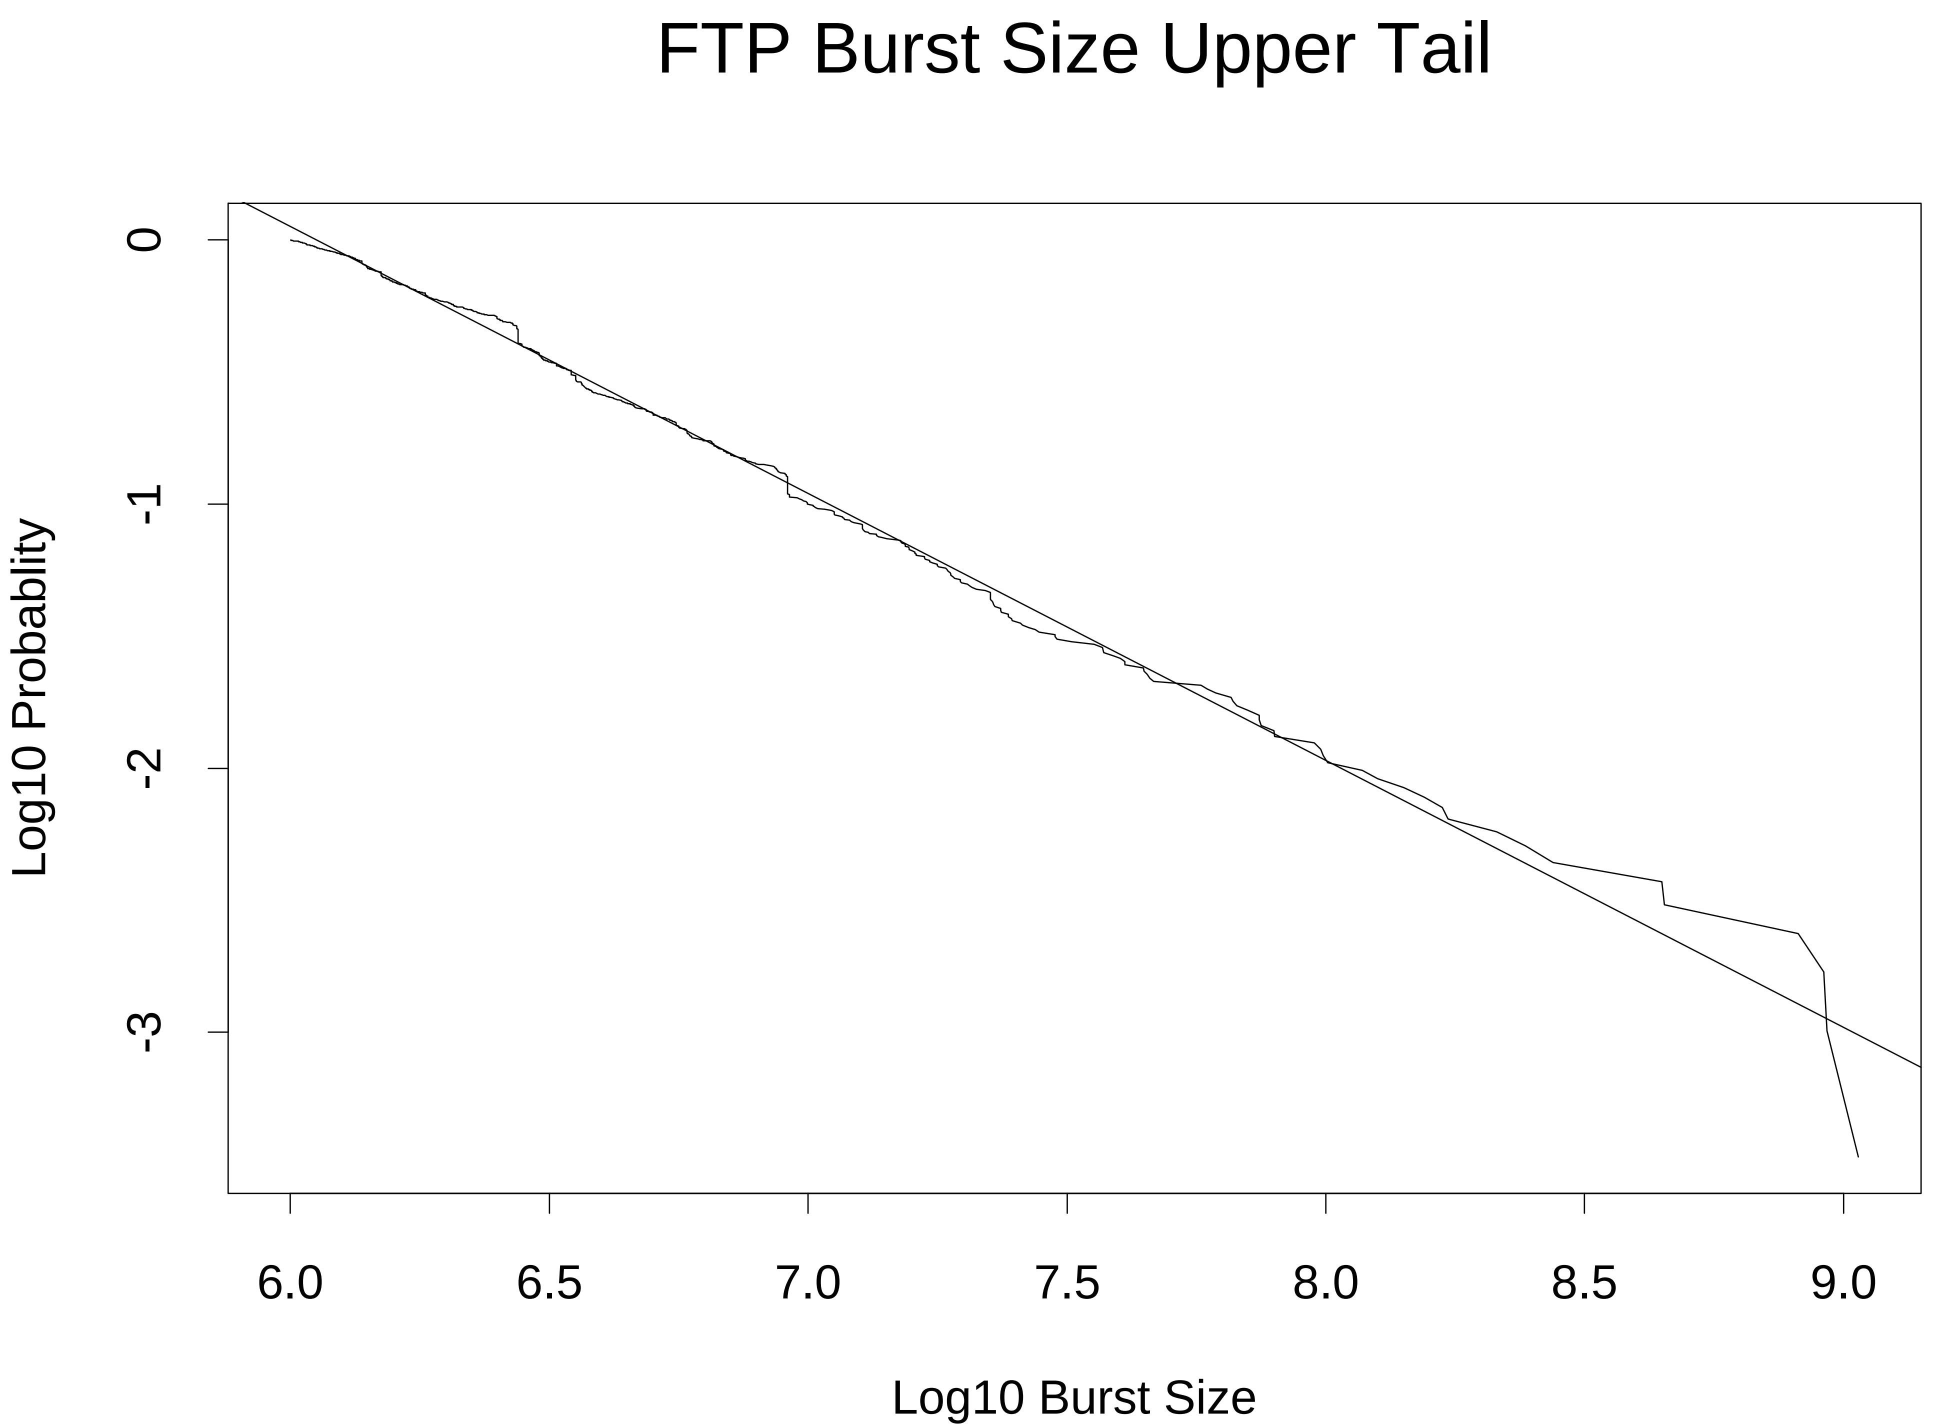

FTP Burst Size Upper Tail
Log10 Probablity
0
-1
-2
-3
6.0
6.5
7.0
7.5
8.0
8.5
9.0
Log10 Burst Size

Testing Poisson models:

How do you test whether Poisson models are appropriate?

- Should see *exponentially distributed* interarrivals.
- Should see *independent* interarrivals.

What λ should you test for?

How do you test for exponentially distributed interarrivals?

- statistical distribution tests (e.g., χ^2)
- graphical tests

How do you test for independent interarrivals?

- statistical tests for *correlation*

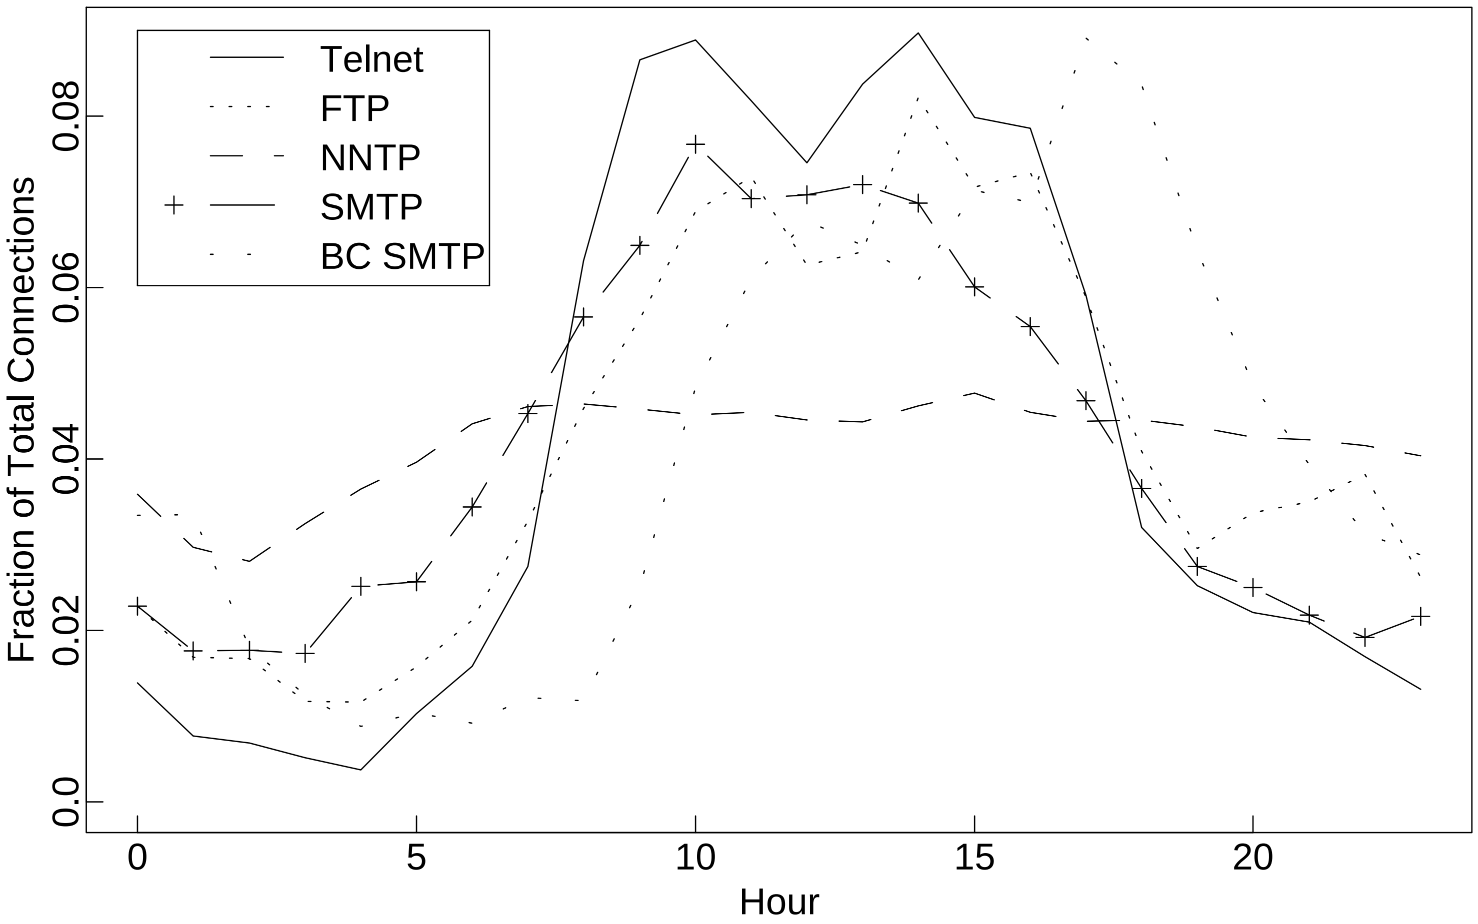
Telnet
FTP
NNTP
SMTP
BC SMTP
Fraction of Total Connections
0.0
0.02
0.04
0.06
0.08
0
5
10
15
20
Hour

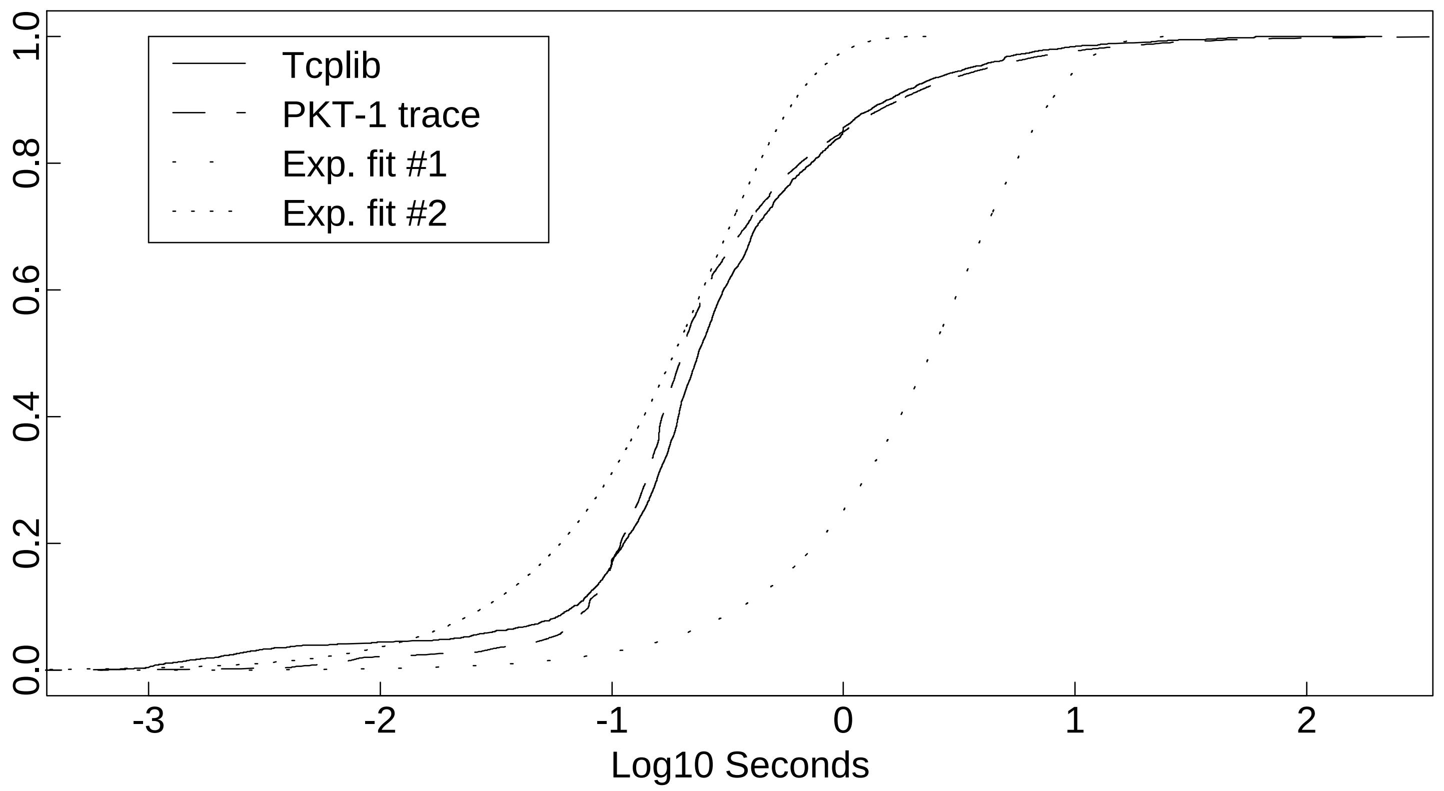
Tcplib
PKT-1 trace
Exp. fit #1
Exp. fit #2
0.0
0.2
0.4
0.6
0.8
1.0
-3
-2
-1
0
1
2
Log10 Seconds

Results of tests:

Packet interarrivals *not* exponentially distributed:

[JR86], [Gusella90], [FL91], [DJCME92]

Connection interarrivals *not* exponentially distributed [PF94]:

- except for user *sessions*, when viewed over 1-hour intervals.
- FTP, NNTP, SMTP also show significant correlations.

The failure of Poisson modeling:

- packet arrivals are not Poisson
- most connection arrivals are not Poisson
- connection sizes are log-normal, not exponential
- packet arrivals are correlated
- connection arrivals are correlated

Poisson modeling fails on all fronts. With what do we replace it?

The fractal revolution:

Poisson models have a *fundamental* problem: they grievously underestimate burstiness.

Consider the *arrival process*: X_k gives the number of packets arriving during interval k (for some fixed interval length T).

Example:

- take 1 hour of traffic from DEC's Internet link (4.3 Mpkt)
- generate Poisson arrivals . . .
- . . . with same mean and variance

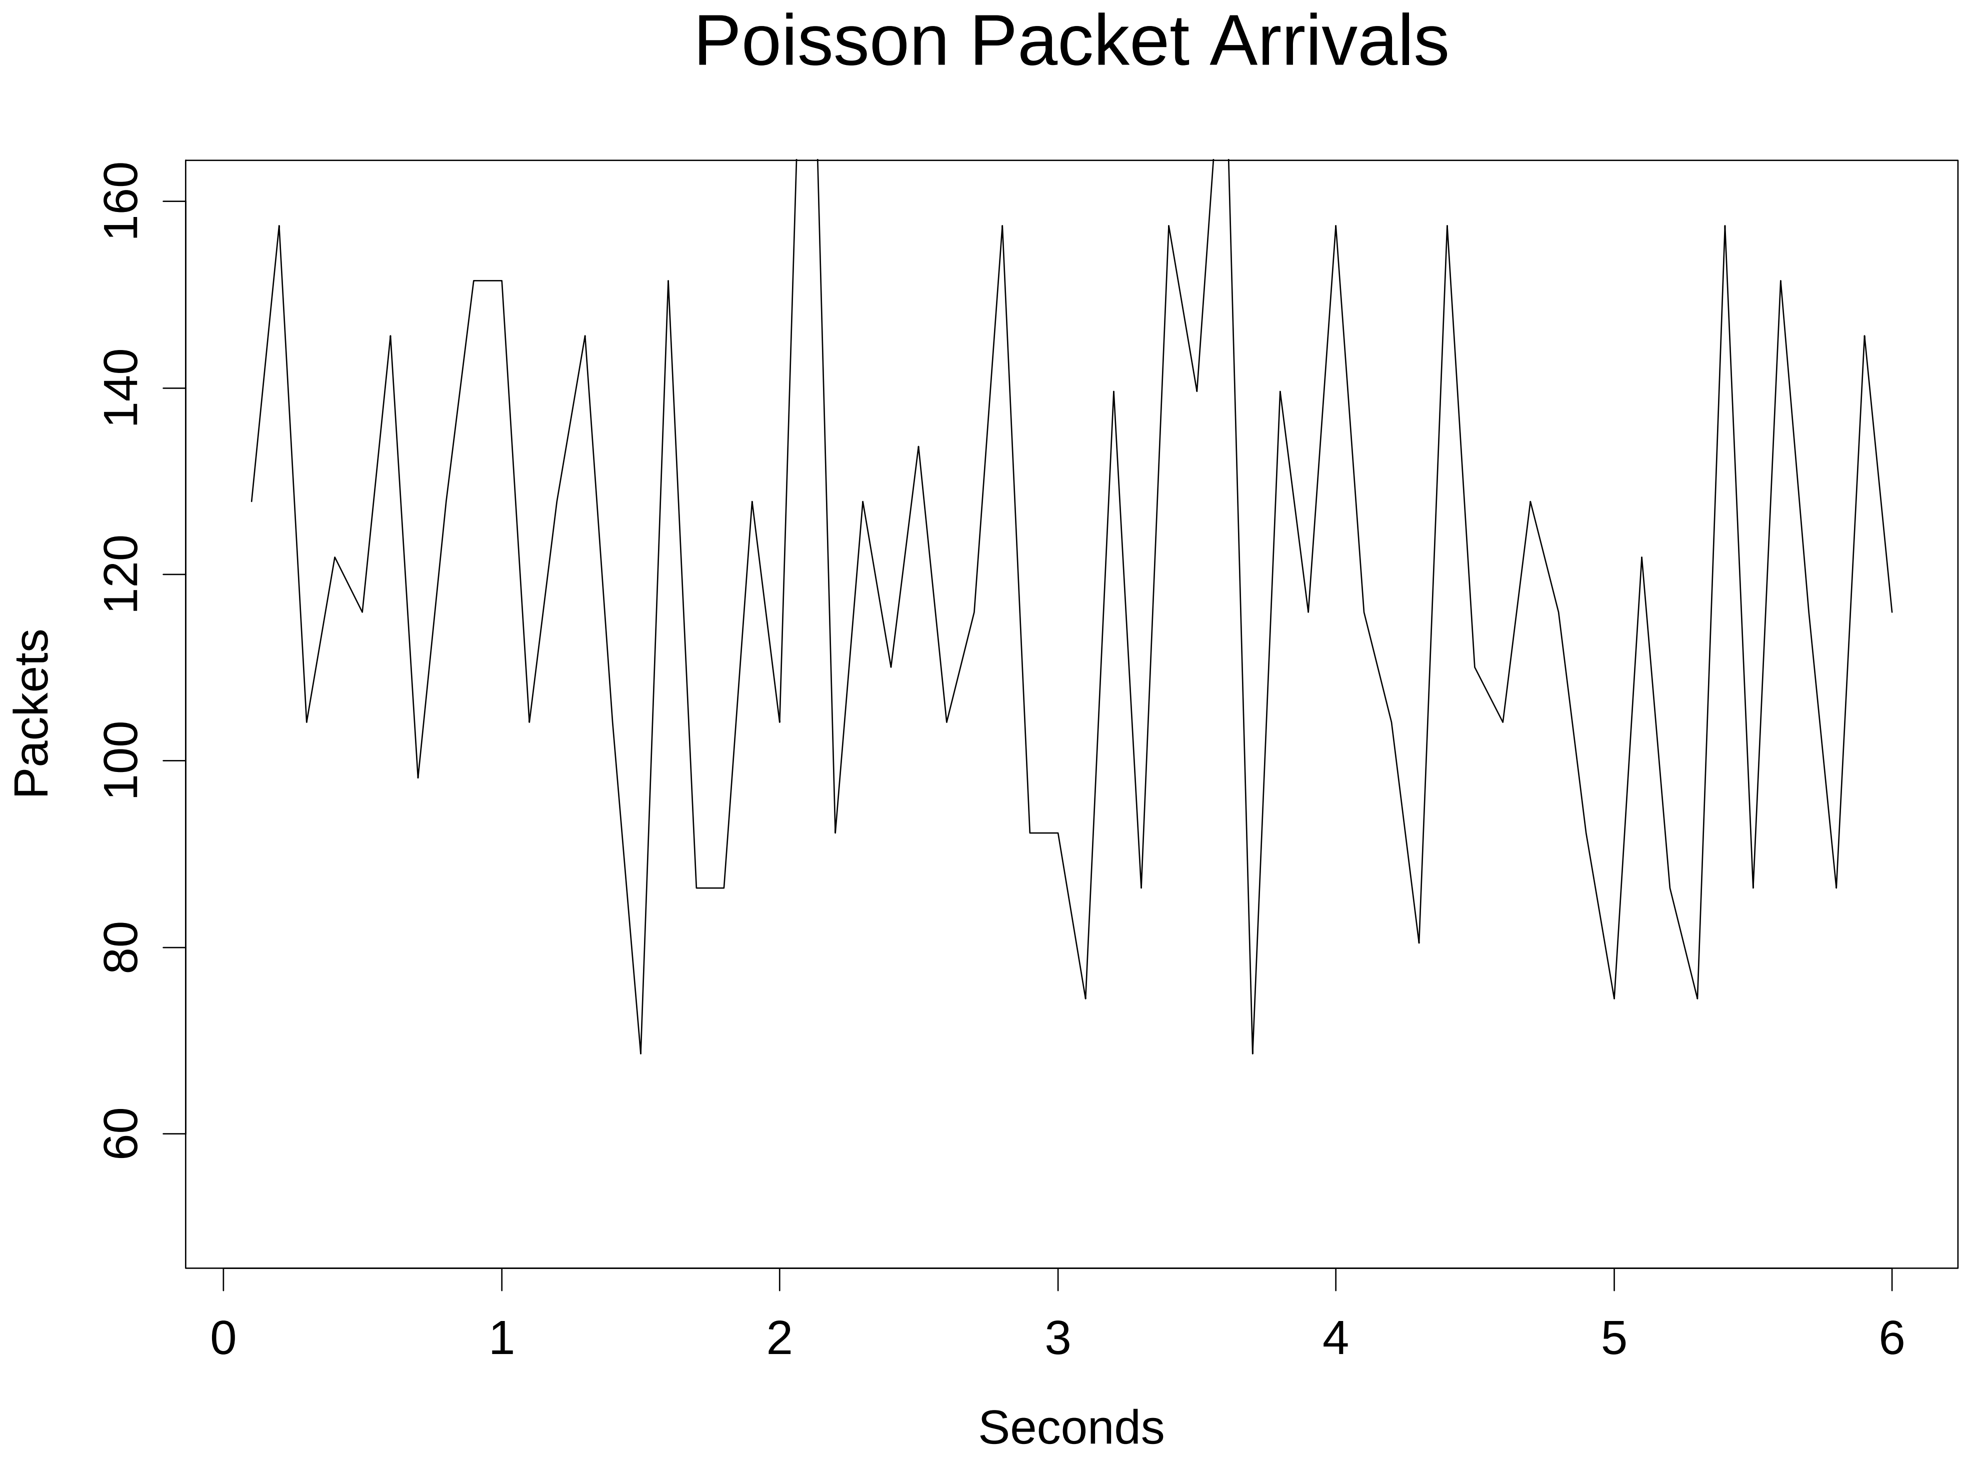

Poisson Packet Arrivals
Packets
160
140
120
100
80
60
0
1
2
3
4
5
6
Seconds

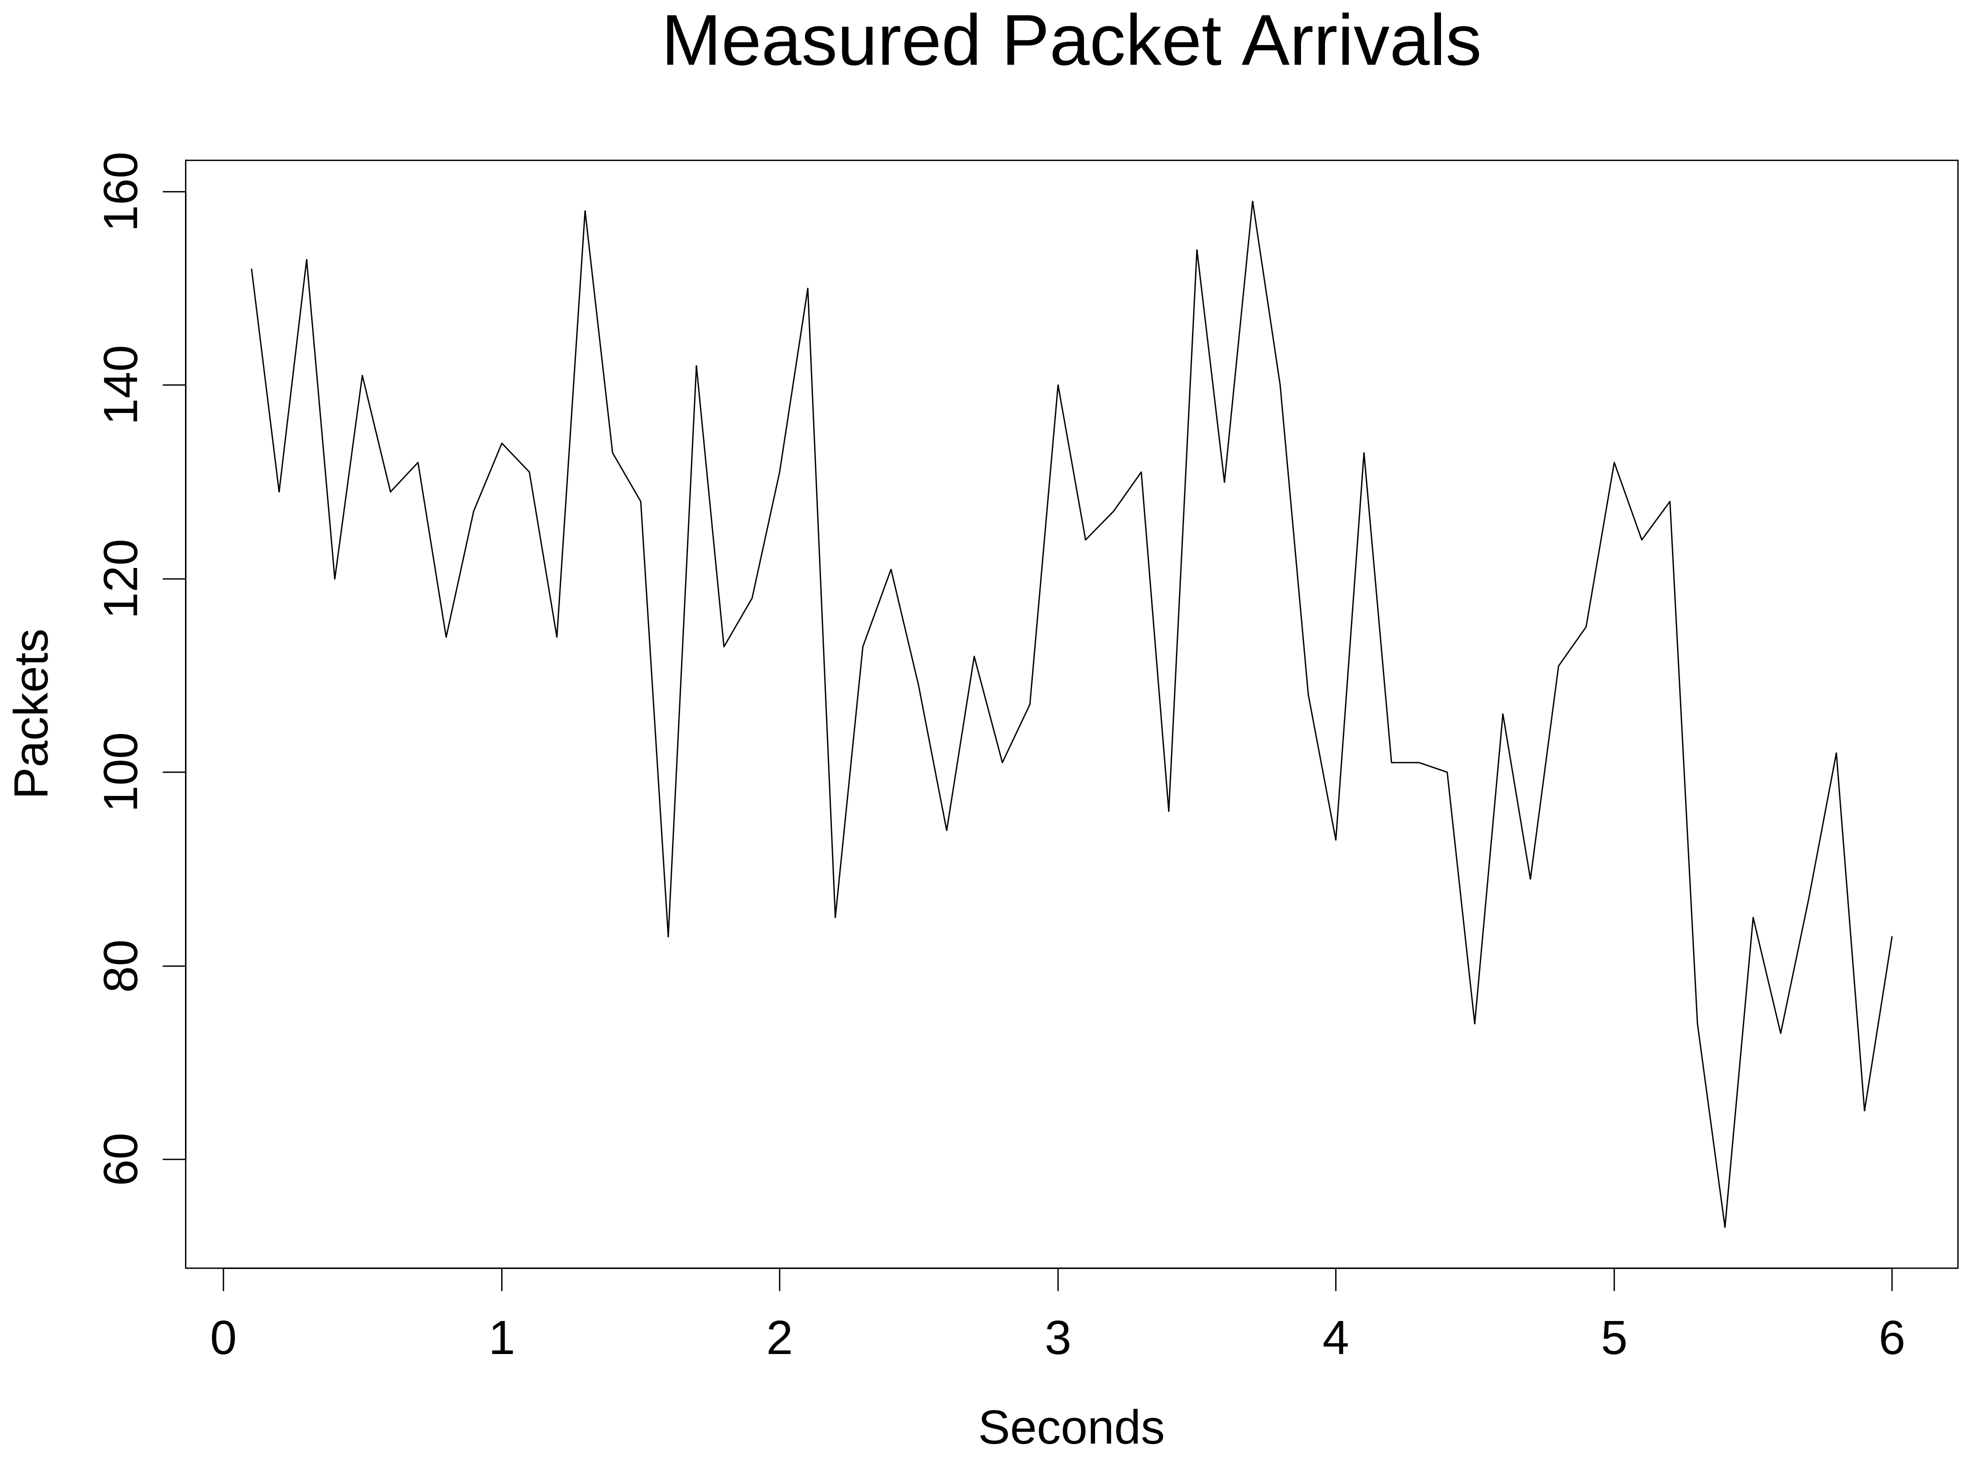
Measured Packet Arrivals
Packets
160
140
120
100
80
60
0
1
2
3
4
5
6
Seconds

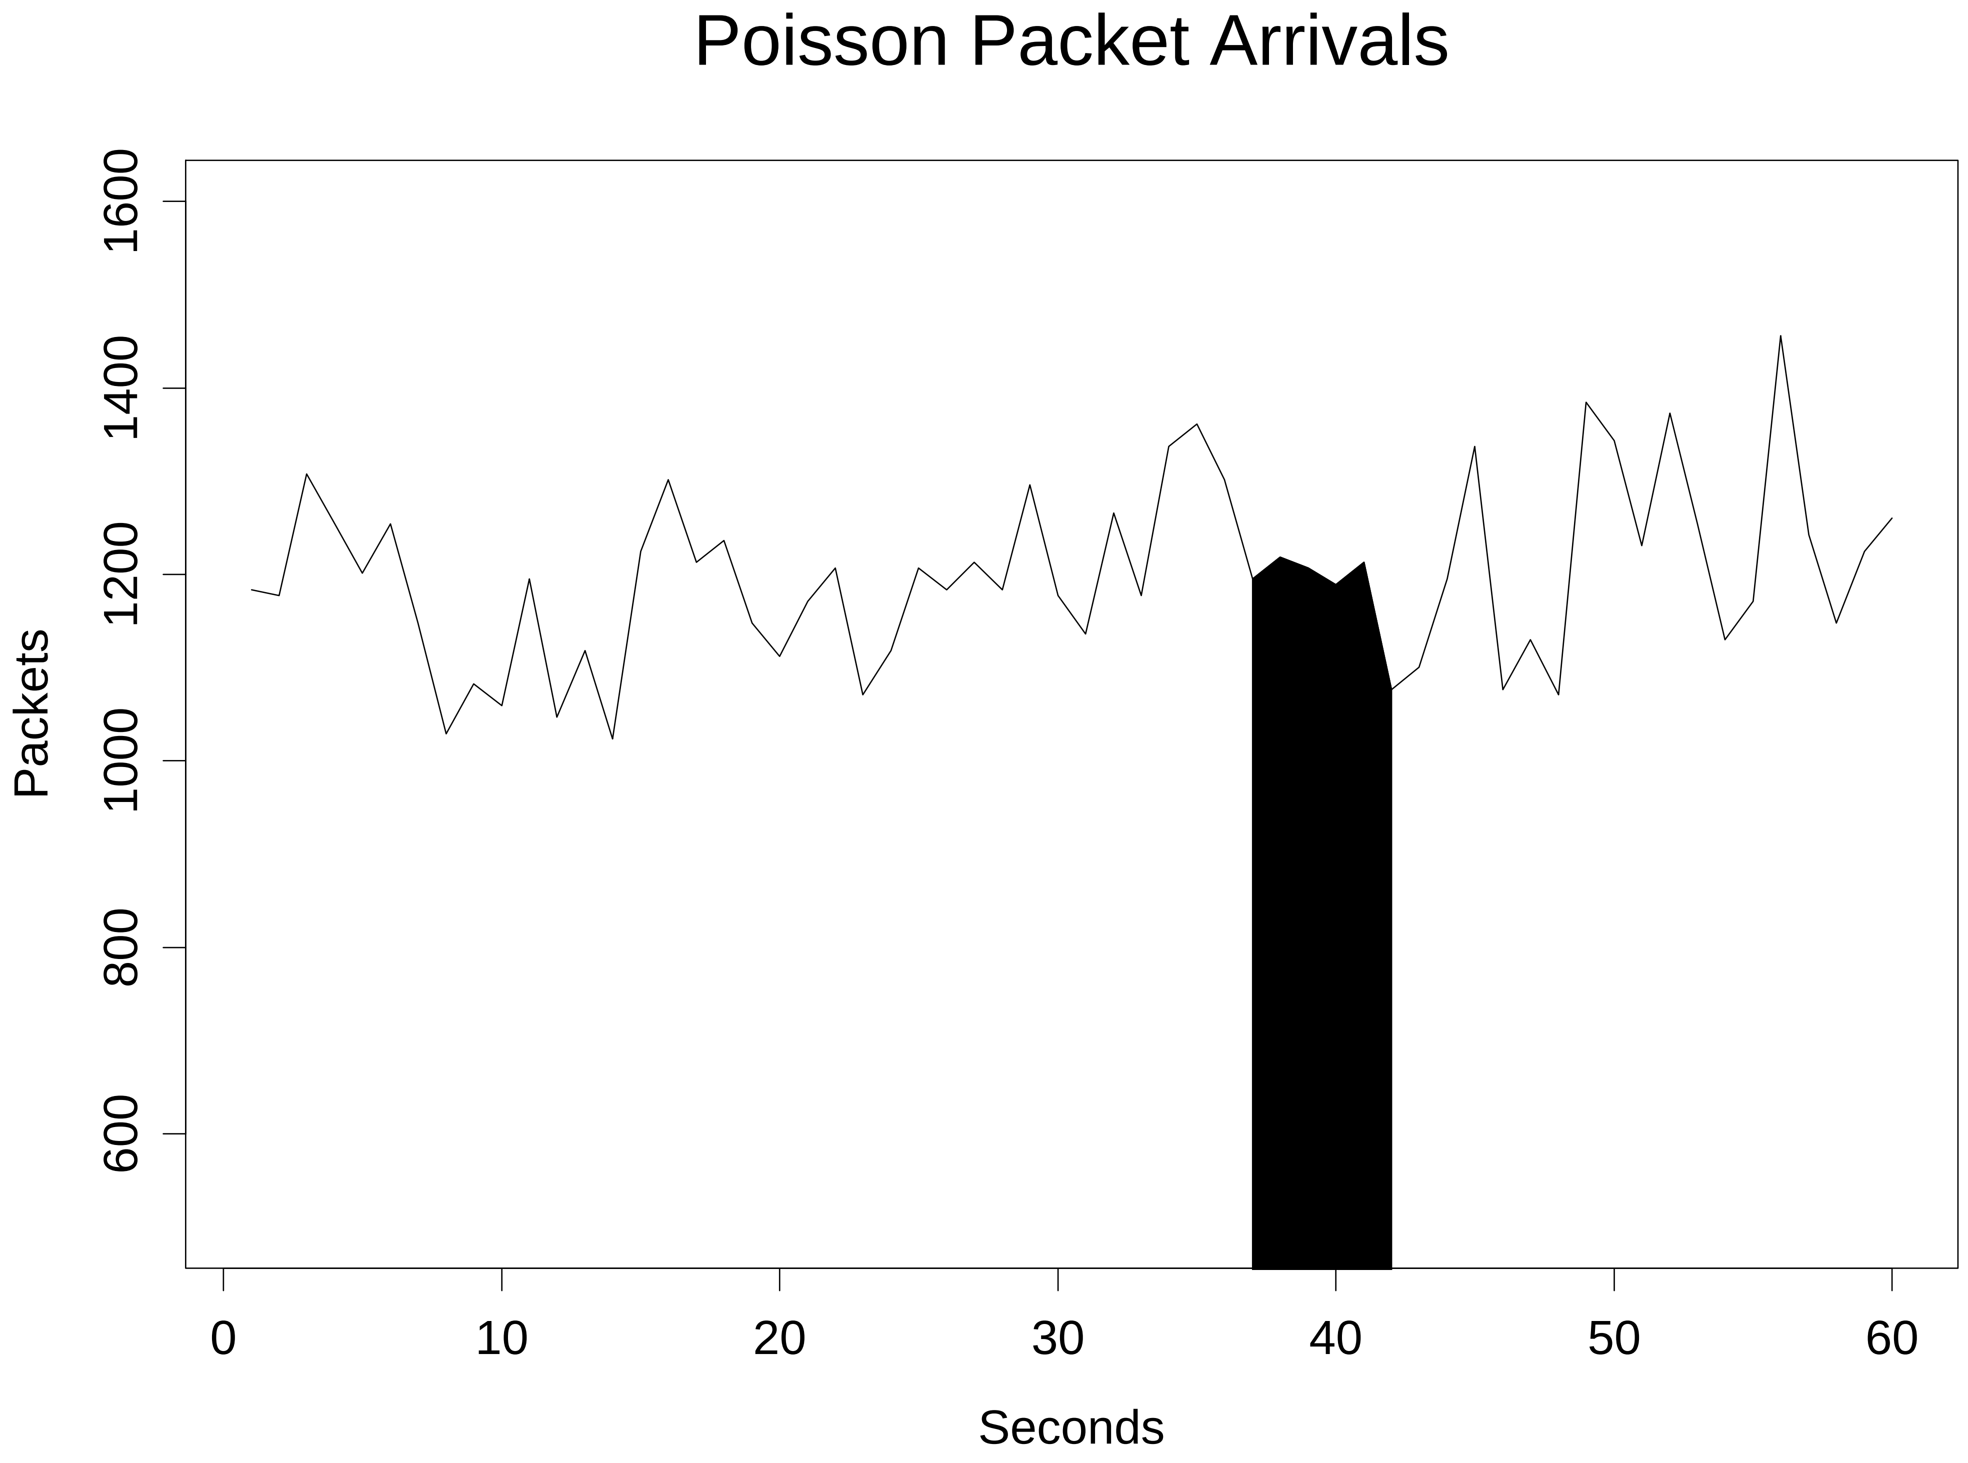
Poisson Packet Arrivals
Packets
600
800
1000
1200
1400
1600
0
10
20
30
40
50
60
Seconds

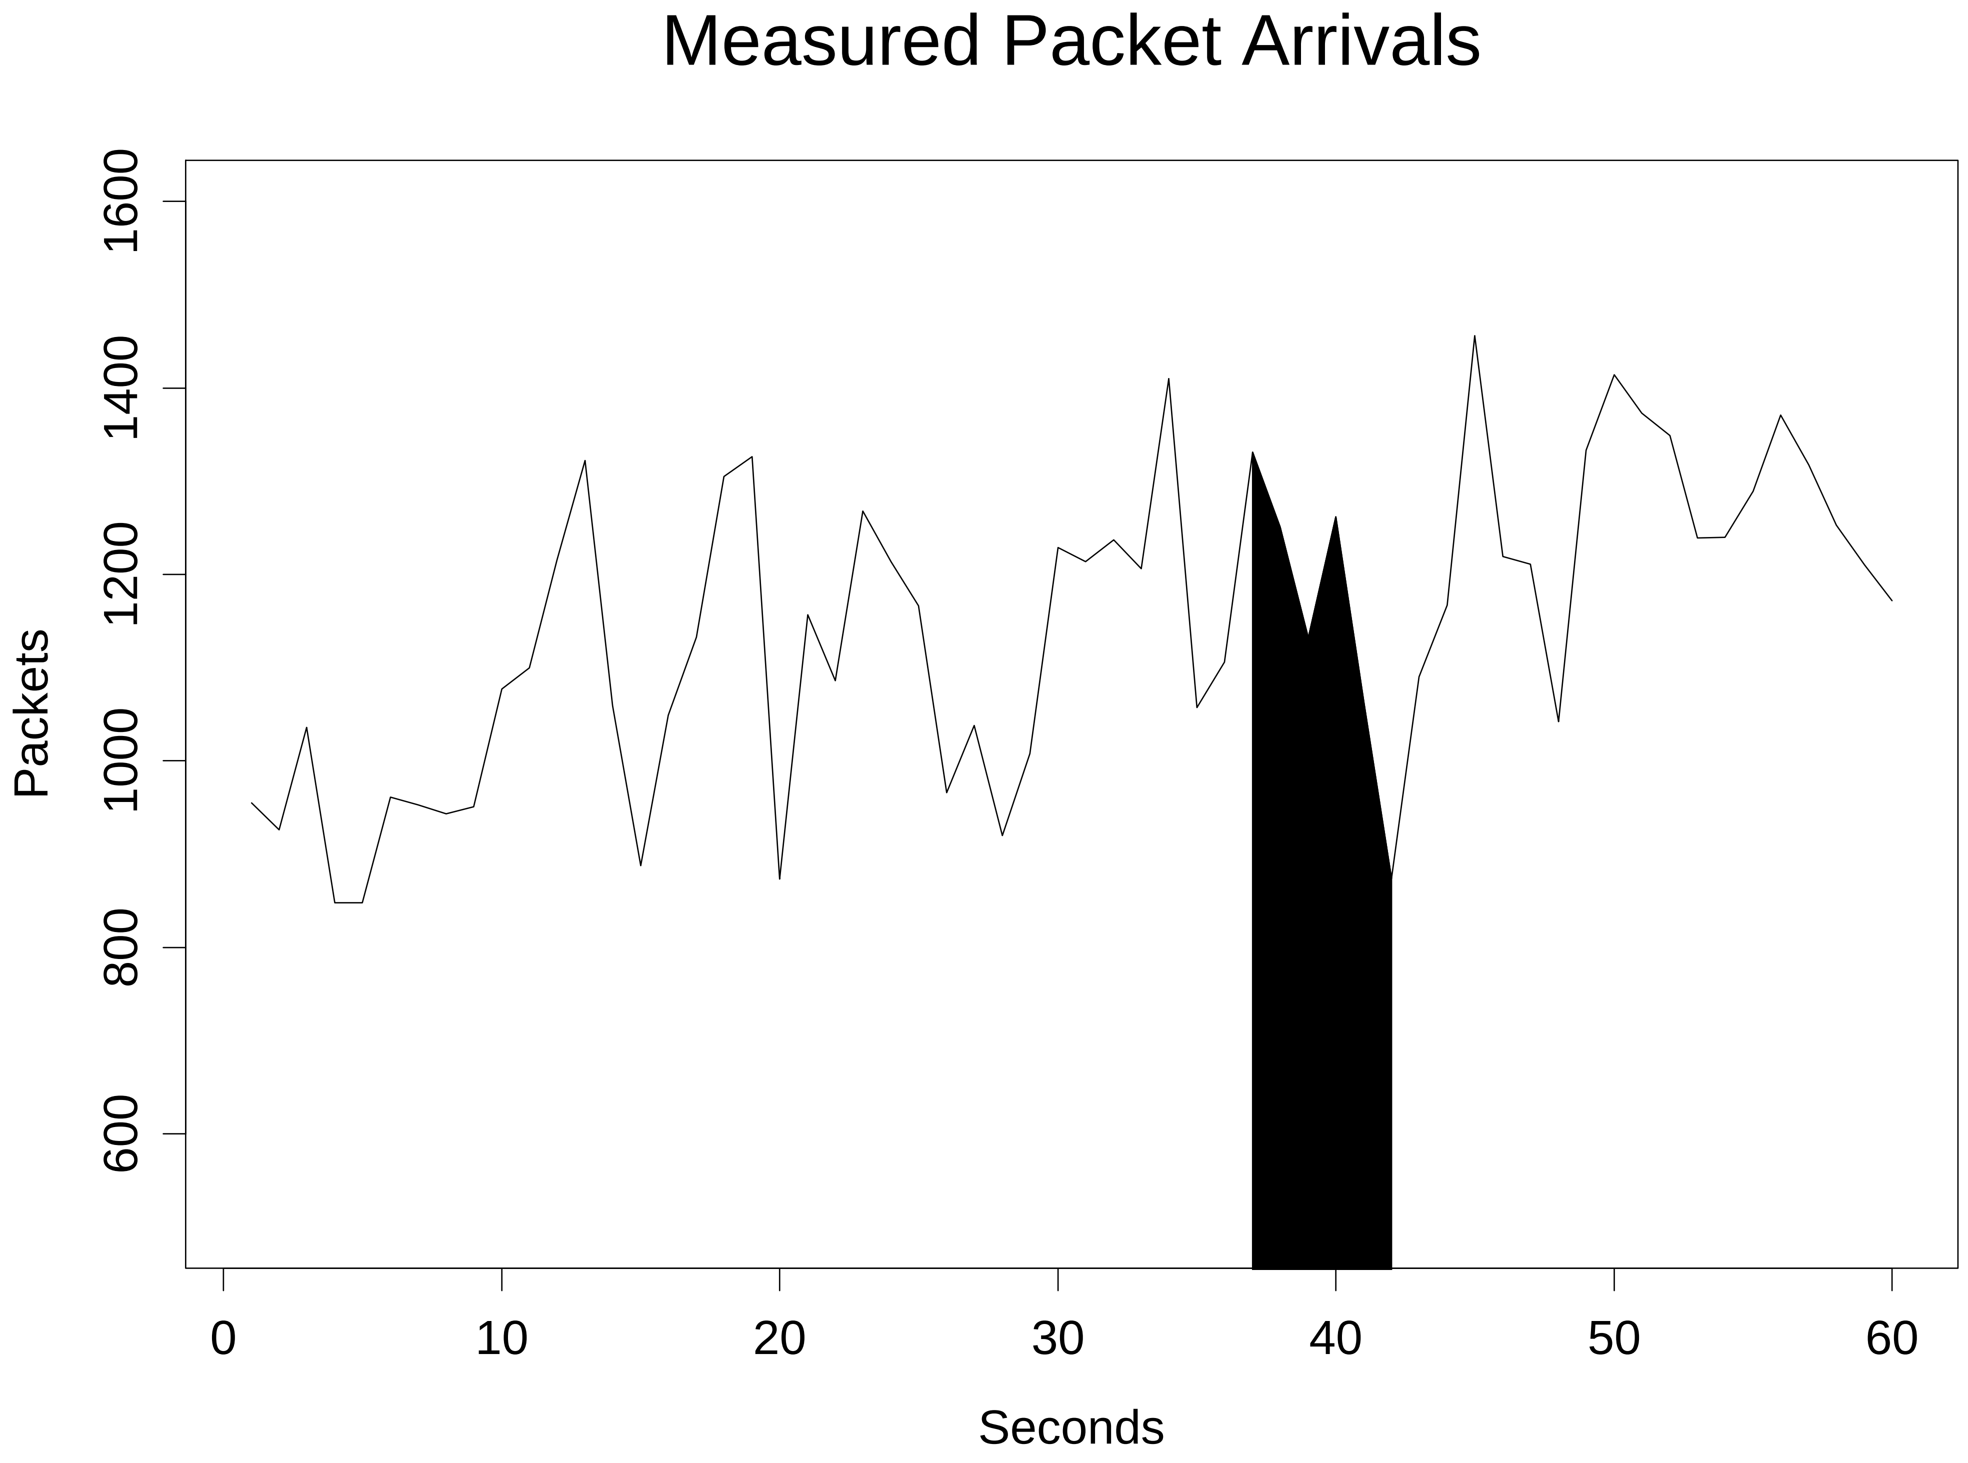

Measured Packet Arrivals
Packets
600
800
1000
1200
1400
1600
0
10
20
30
40
50
60
Seconds

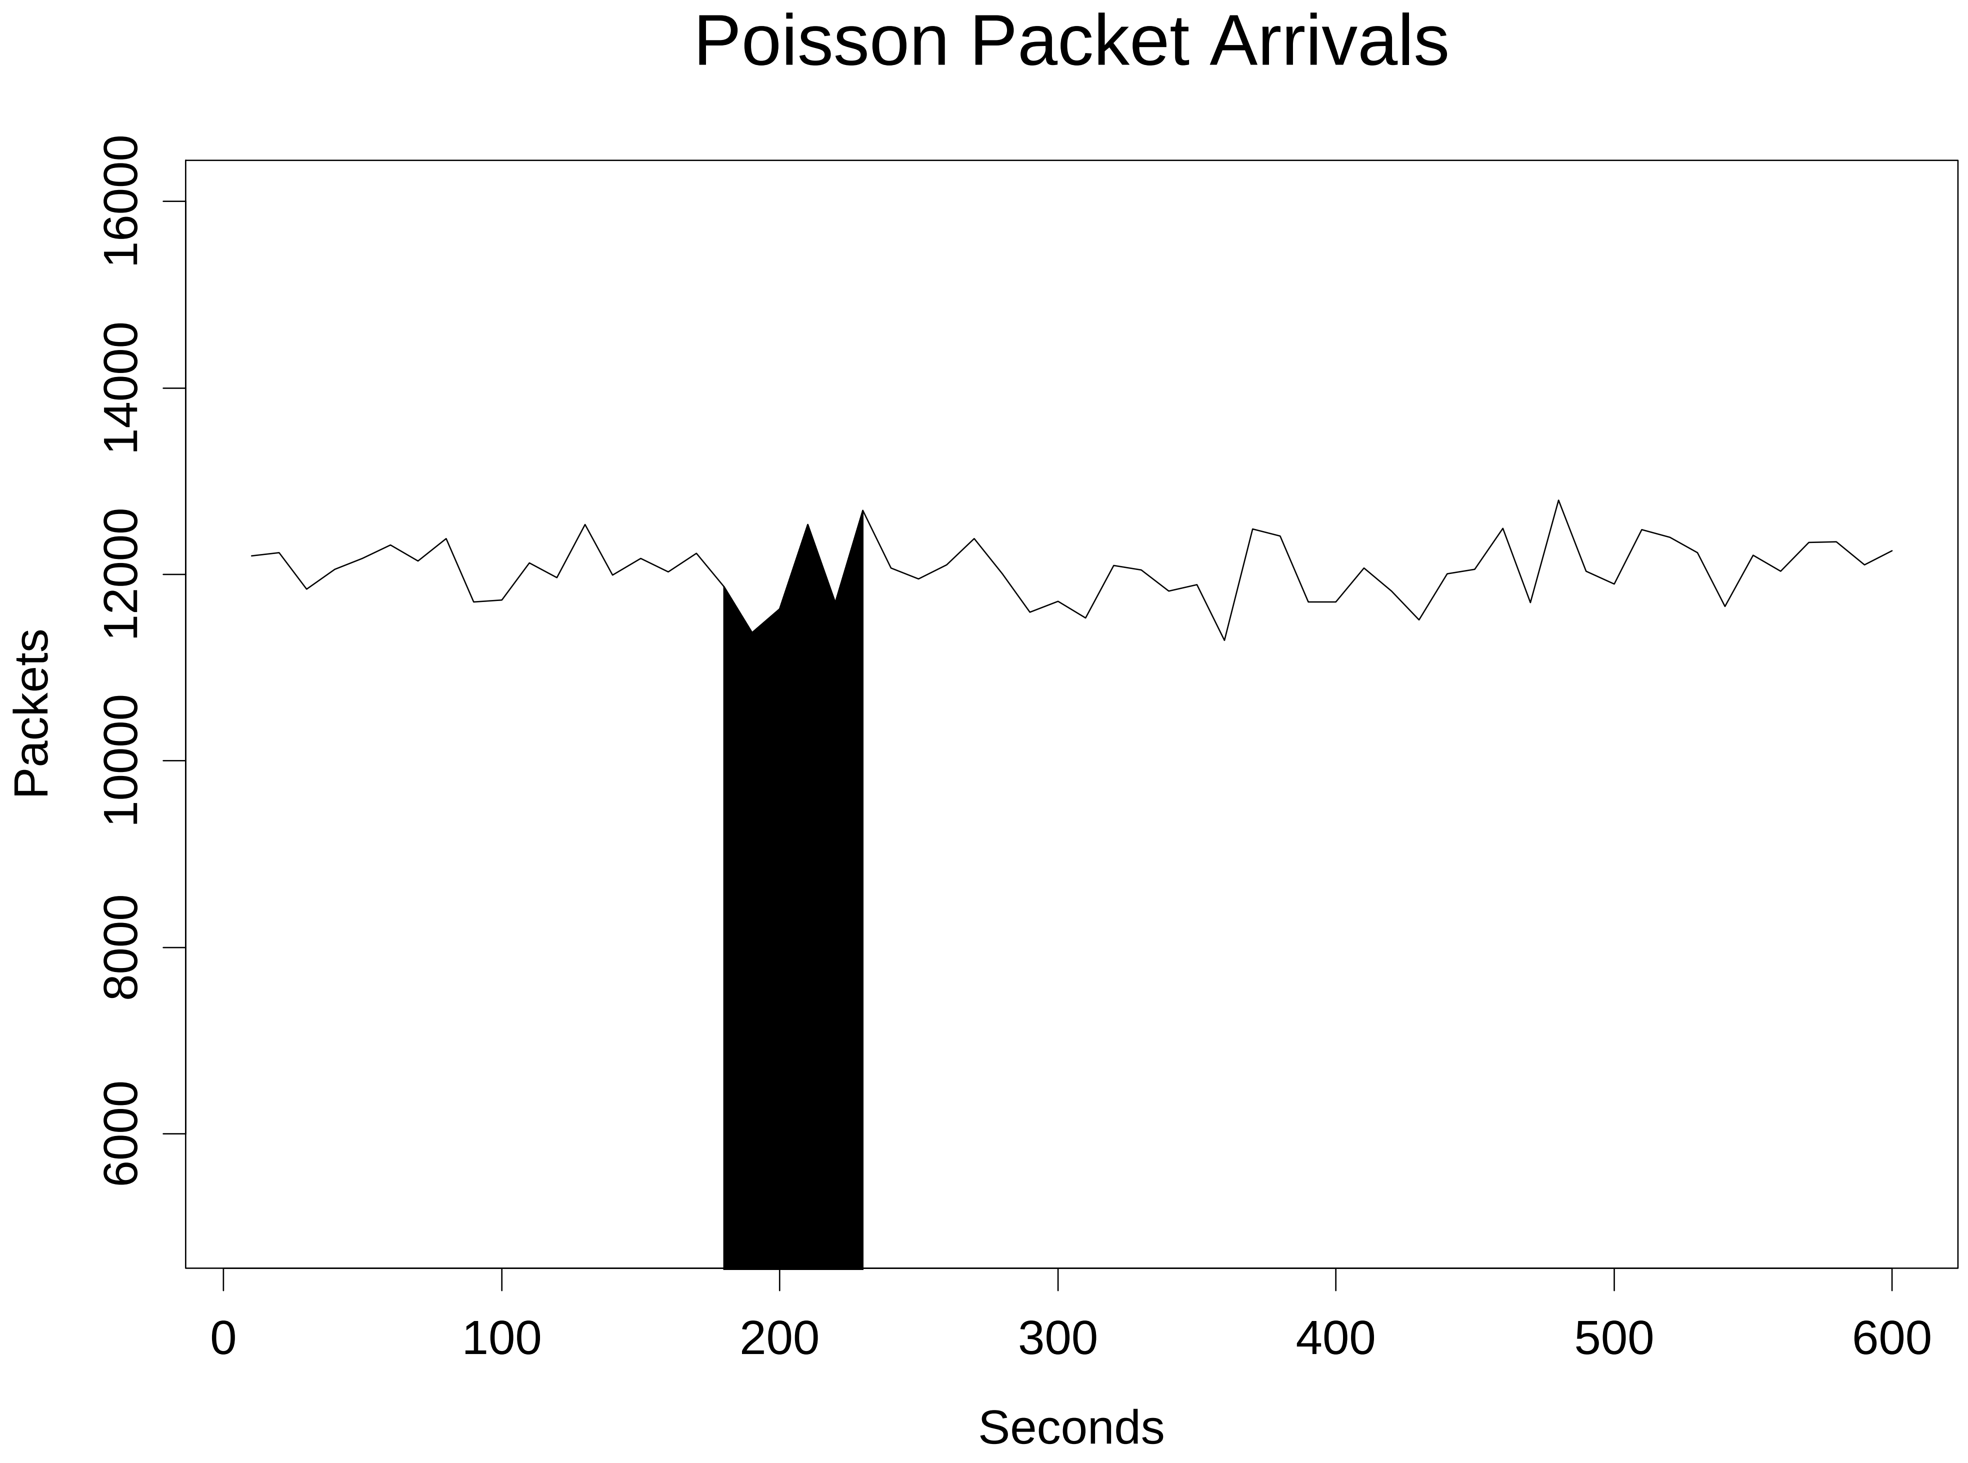
Poisson Packet Arrivals
Packets
6000
8000
10000
12000
14000
16000
Seconds
0
100
200
300
400
500
600

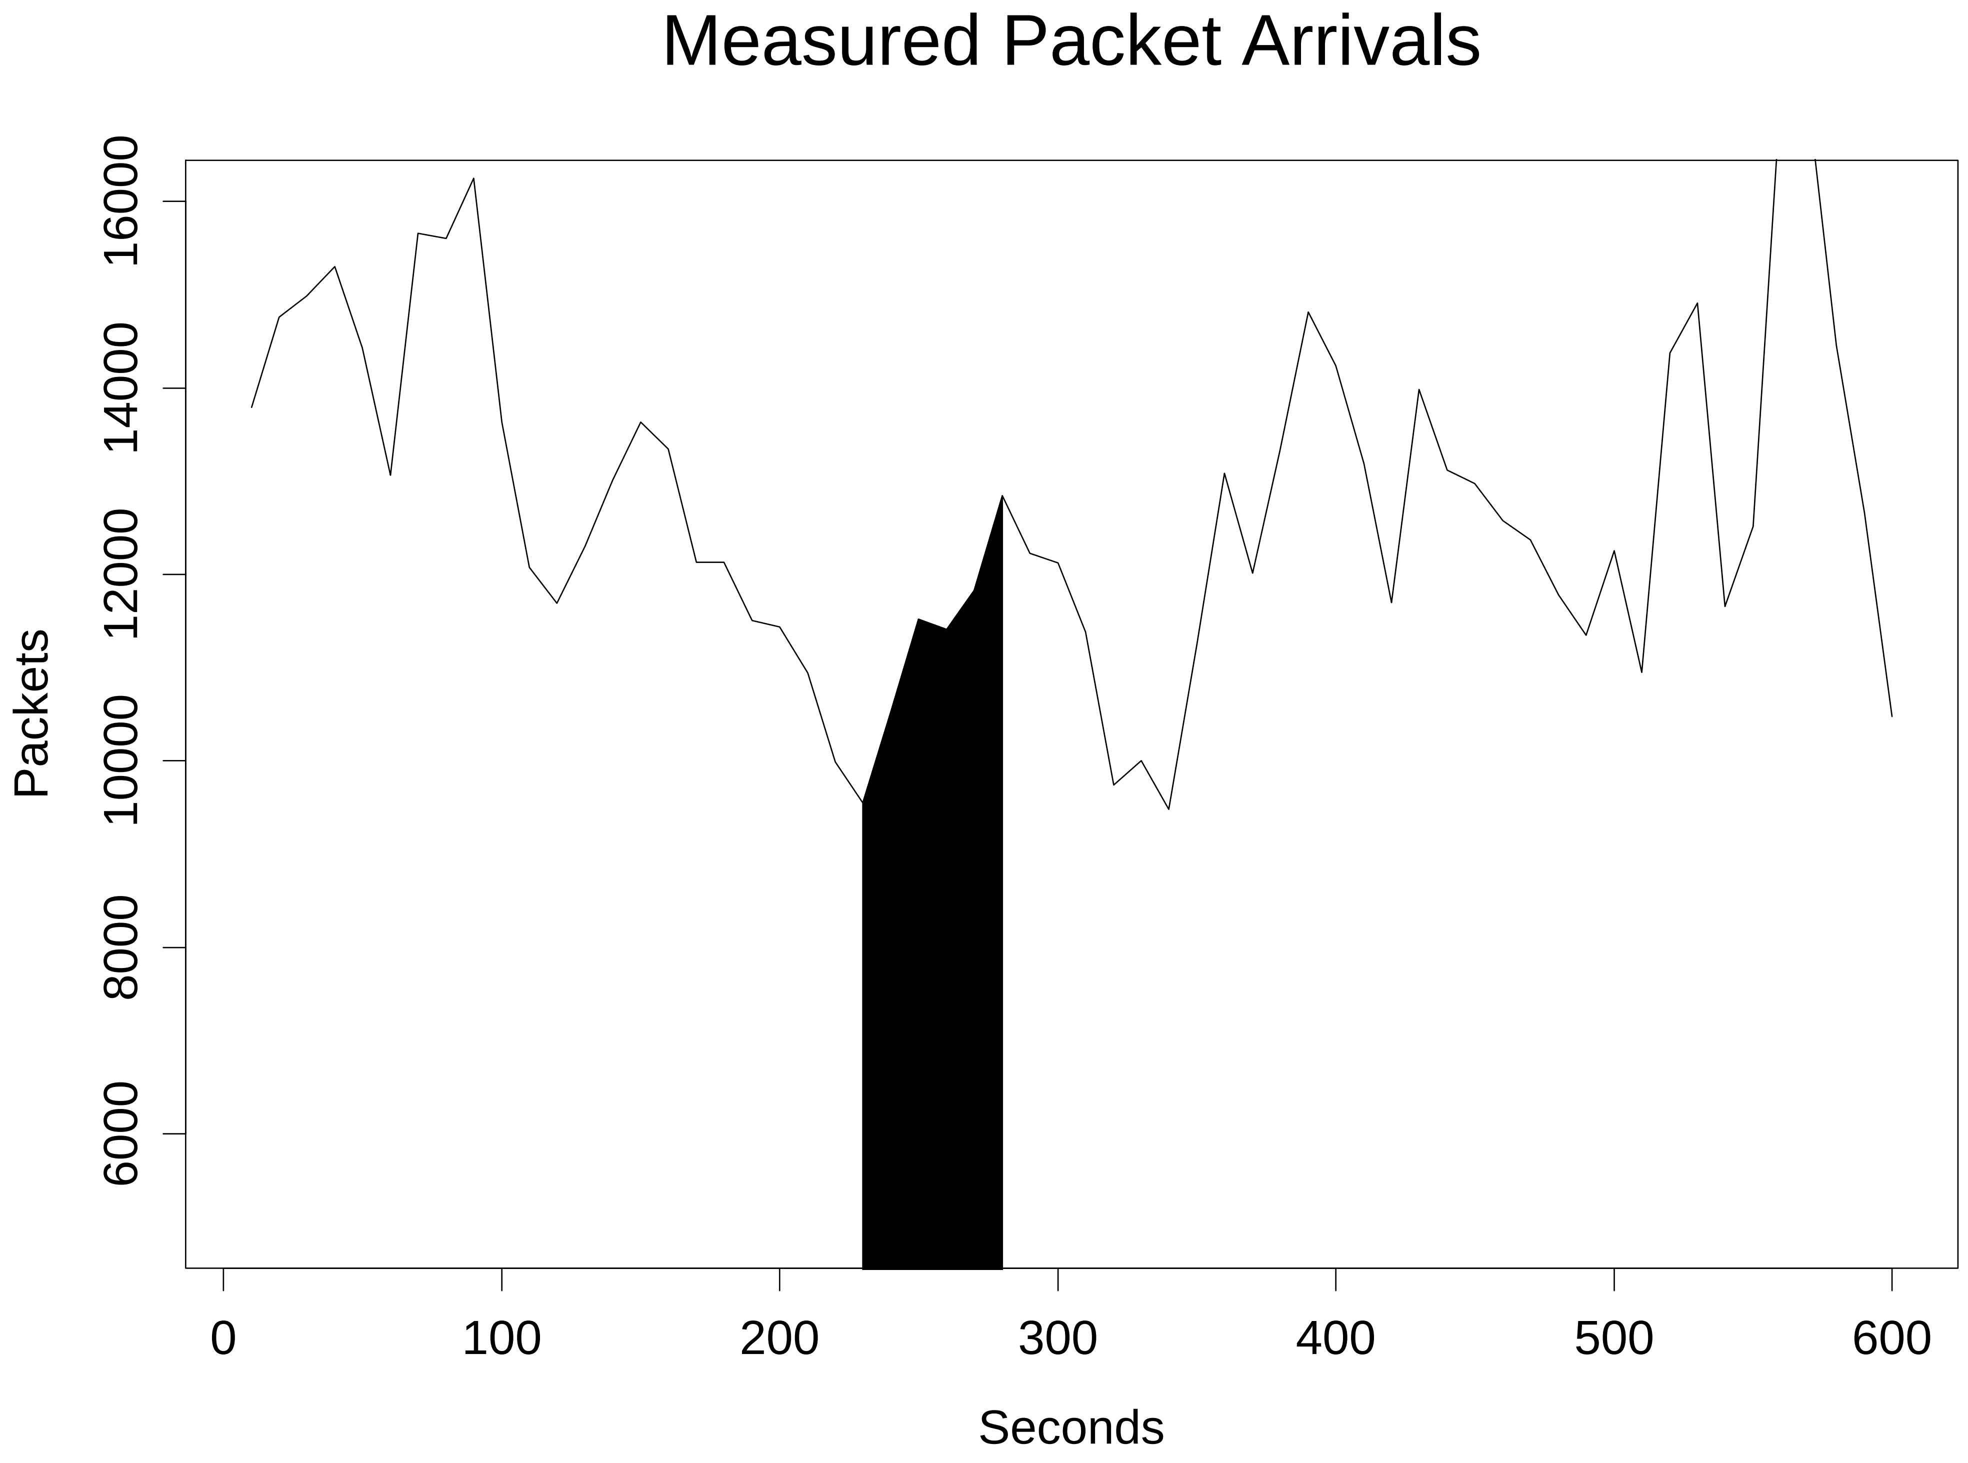
Measured Packet Arrivals
Packets
6000
8000
10000
12000
14000
16000
0
100
200
300
400
500
600
Seconds

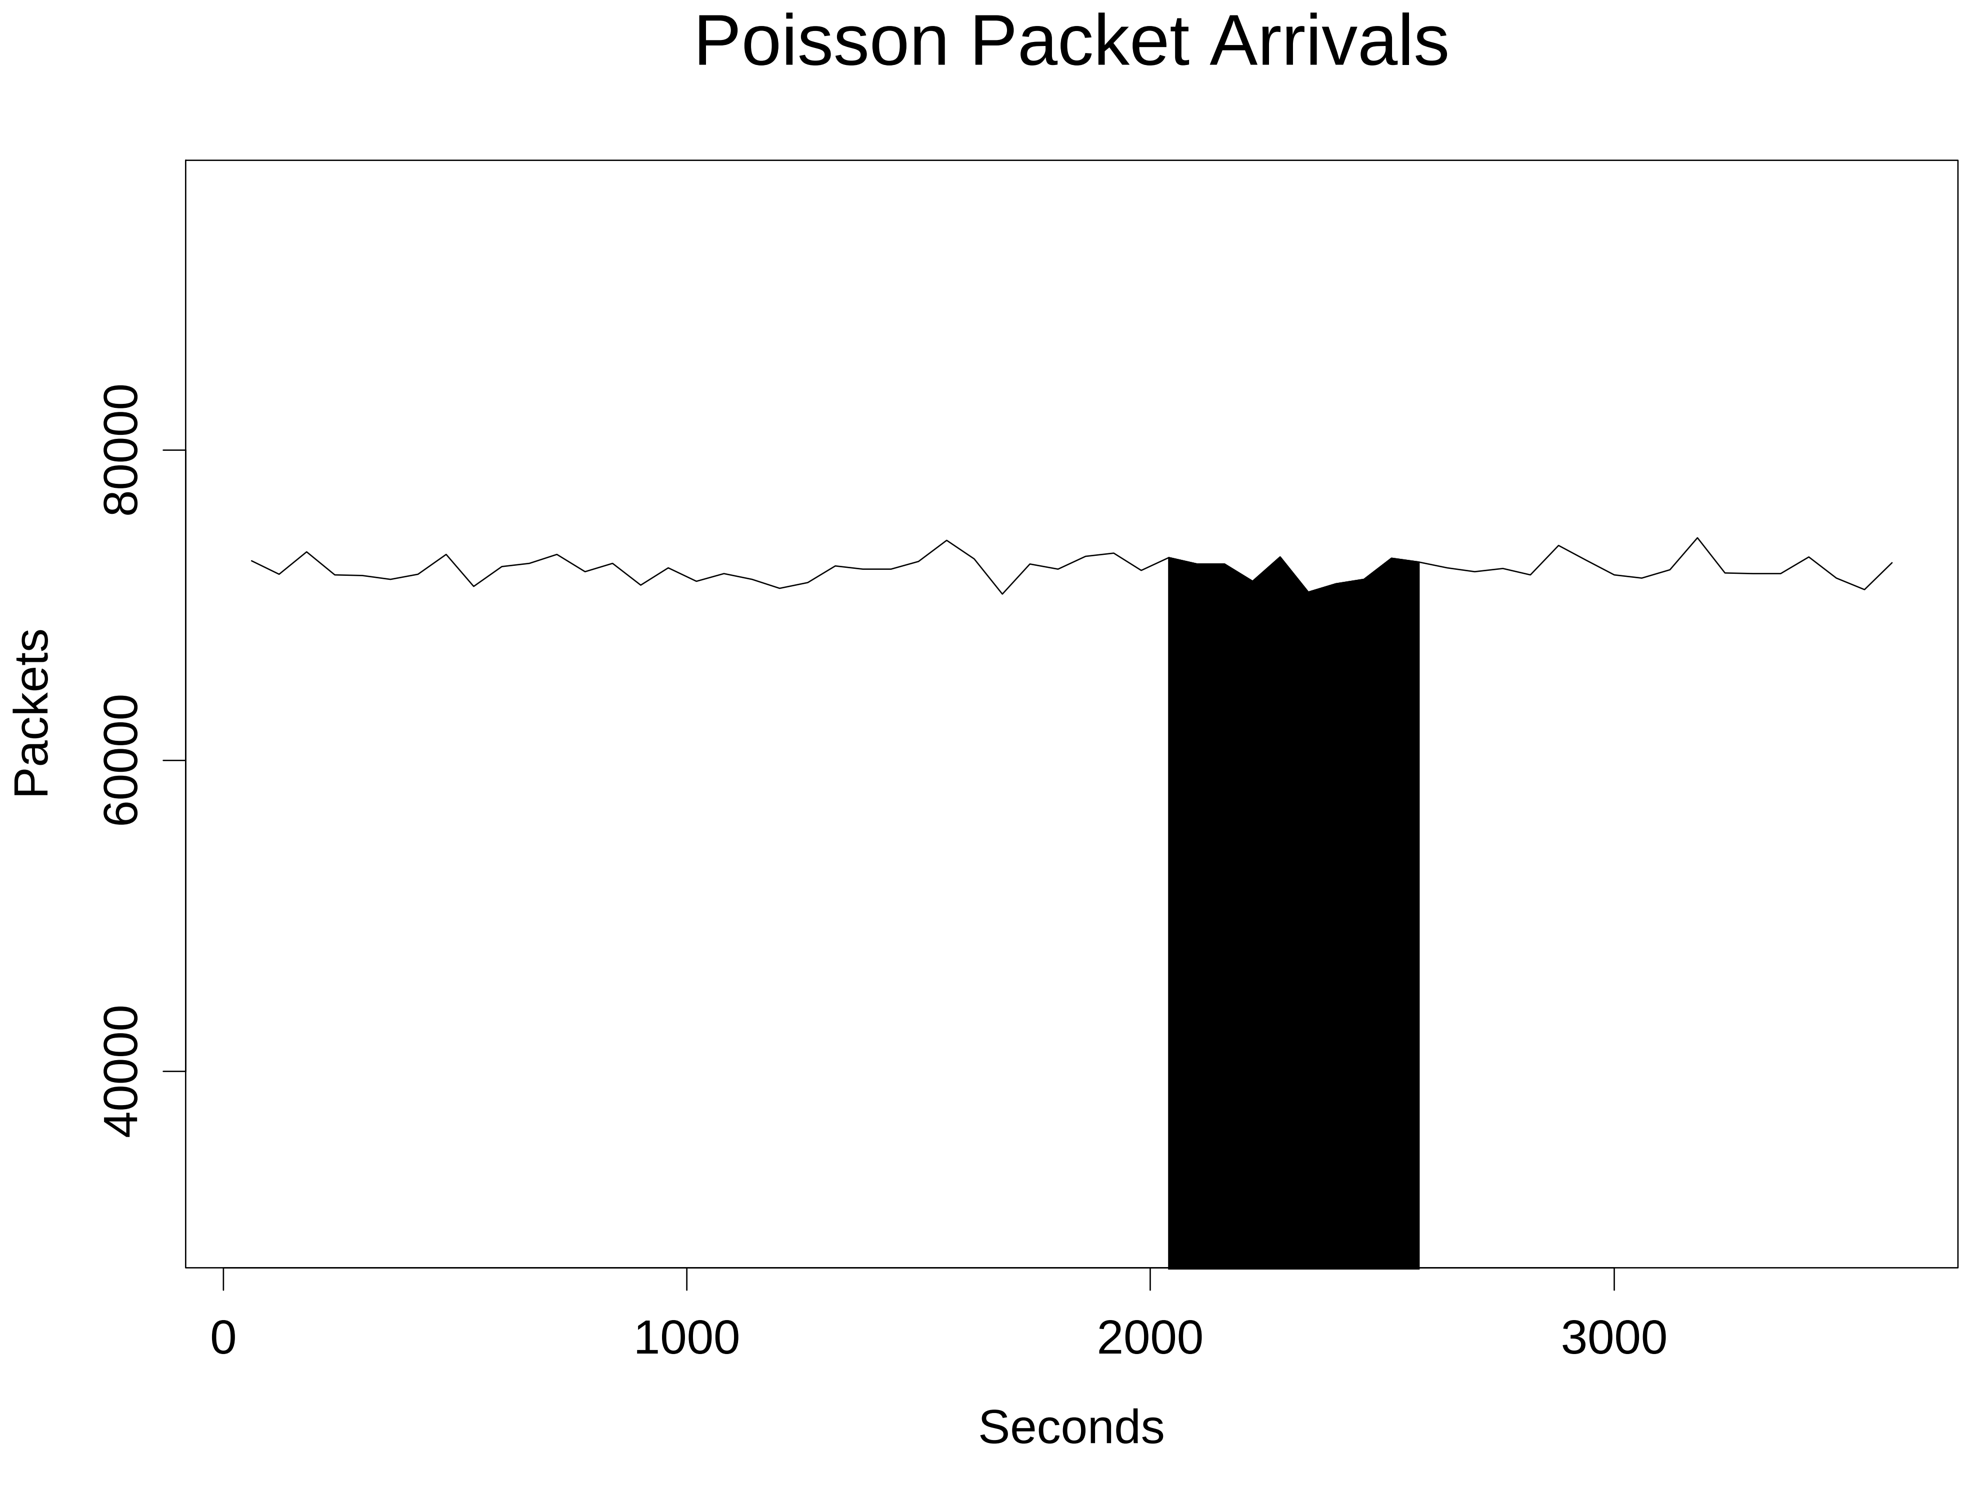
Poisson Packet Arrivals
Packets
80000
60000
40000
0
1000
2000
3000
Seconds

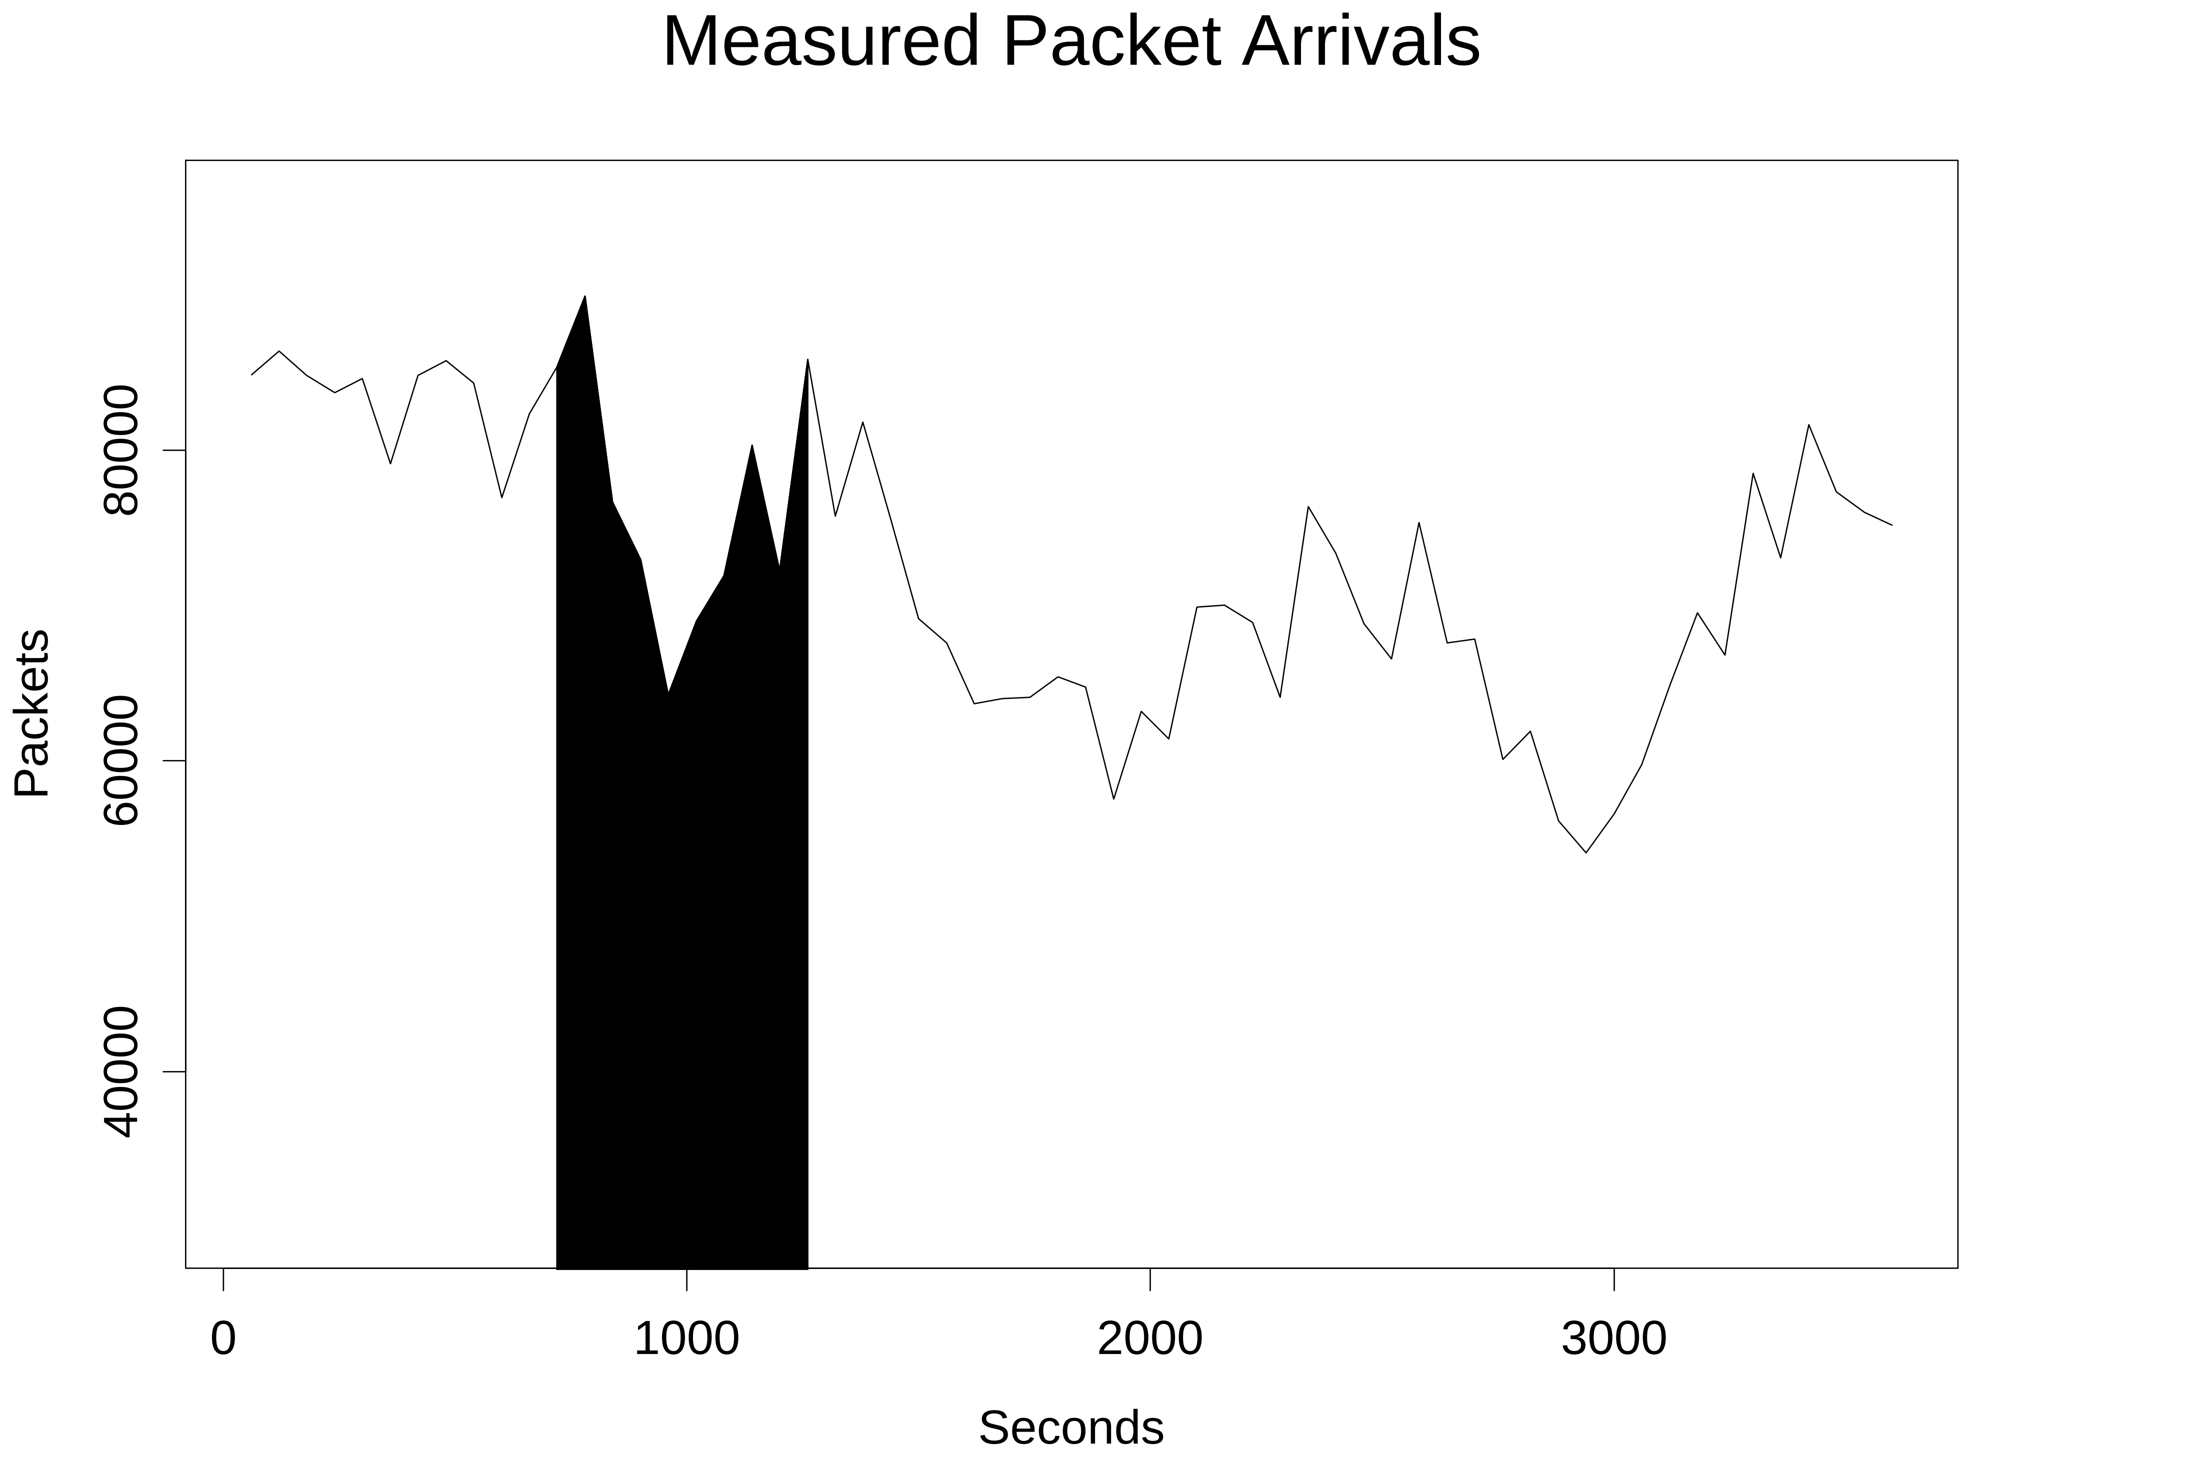
Measured Packet Arrivals
Packets
40000
60000
80000
0
1000
2000
3000
Seconds

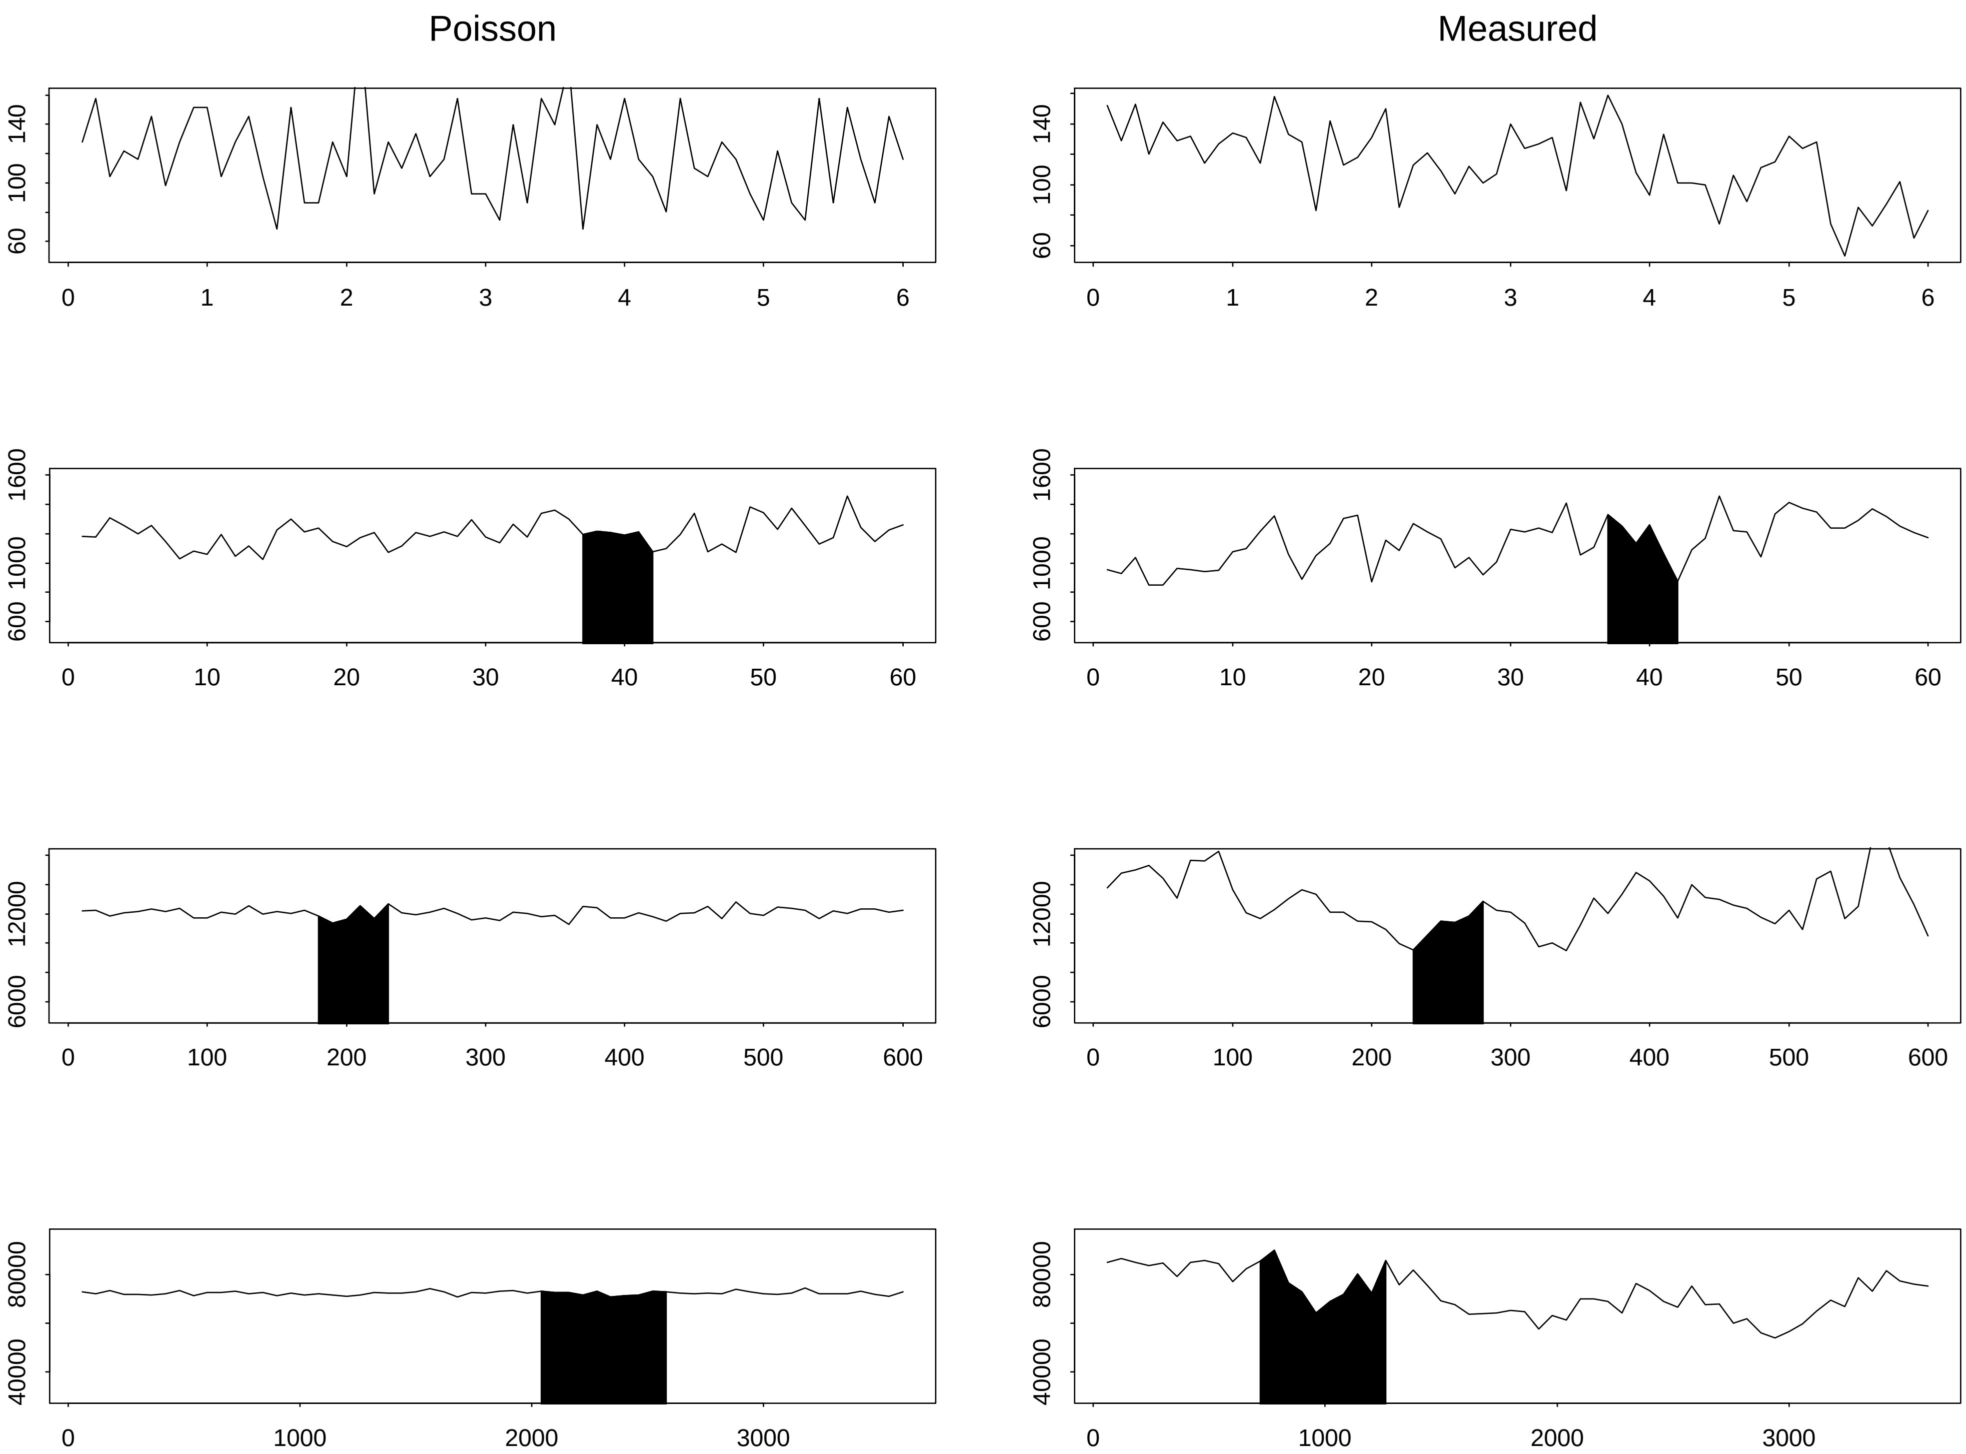

Poisson
Measured
60 100 140
0 1 2 3 4 5 6
600 1000 1600
0 10 20 30 40 50 60
6000 12000
0 100 200 300 400 500 600
40000 80000
0 1000 2000 3000

Burstiness over many time scales:

"Visual self-similarity," per previous plots.

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).

Variance-time plots:

Suppose we have X_k measured for $T = 100$ msec. Considered a "smoothed" version ${X_k}^{(10)}$, corresponding to viewing X_k on a time interval $T' = 10 \times T = 1$ sec:

$$ {X_k}^{(10)} = \frac{1}{10}(X_{10k-10+1} + \cdots + X_{10k}). $$

More generally, for any degree of smoothing m:

$$ {X_k}^{(m)} = \frac{1}{m}(X_{mk-m+1} + \cdots + X_{mk}). $$

Variance-time plots, con't:

A natural measure of burstiness is *variance*:

$$\mathsf{Var}(Y) = E[(Y - E[Y])^2].$$

How does $\mathsf{Var}({X_k}^{(m)})$ vary with m? For a Poisson process:

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

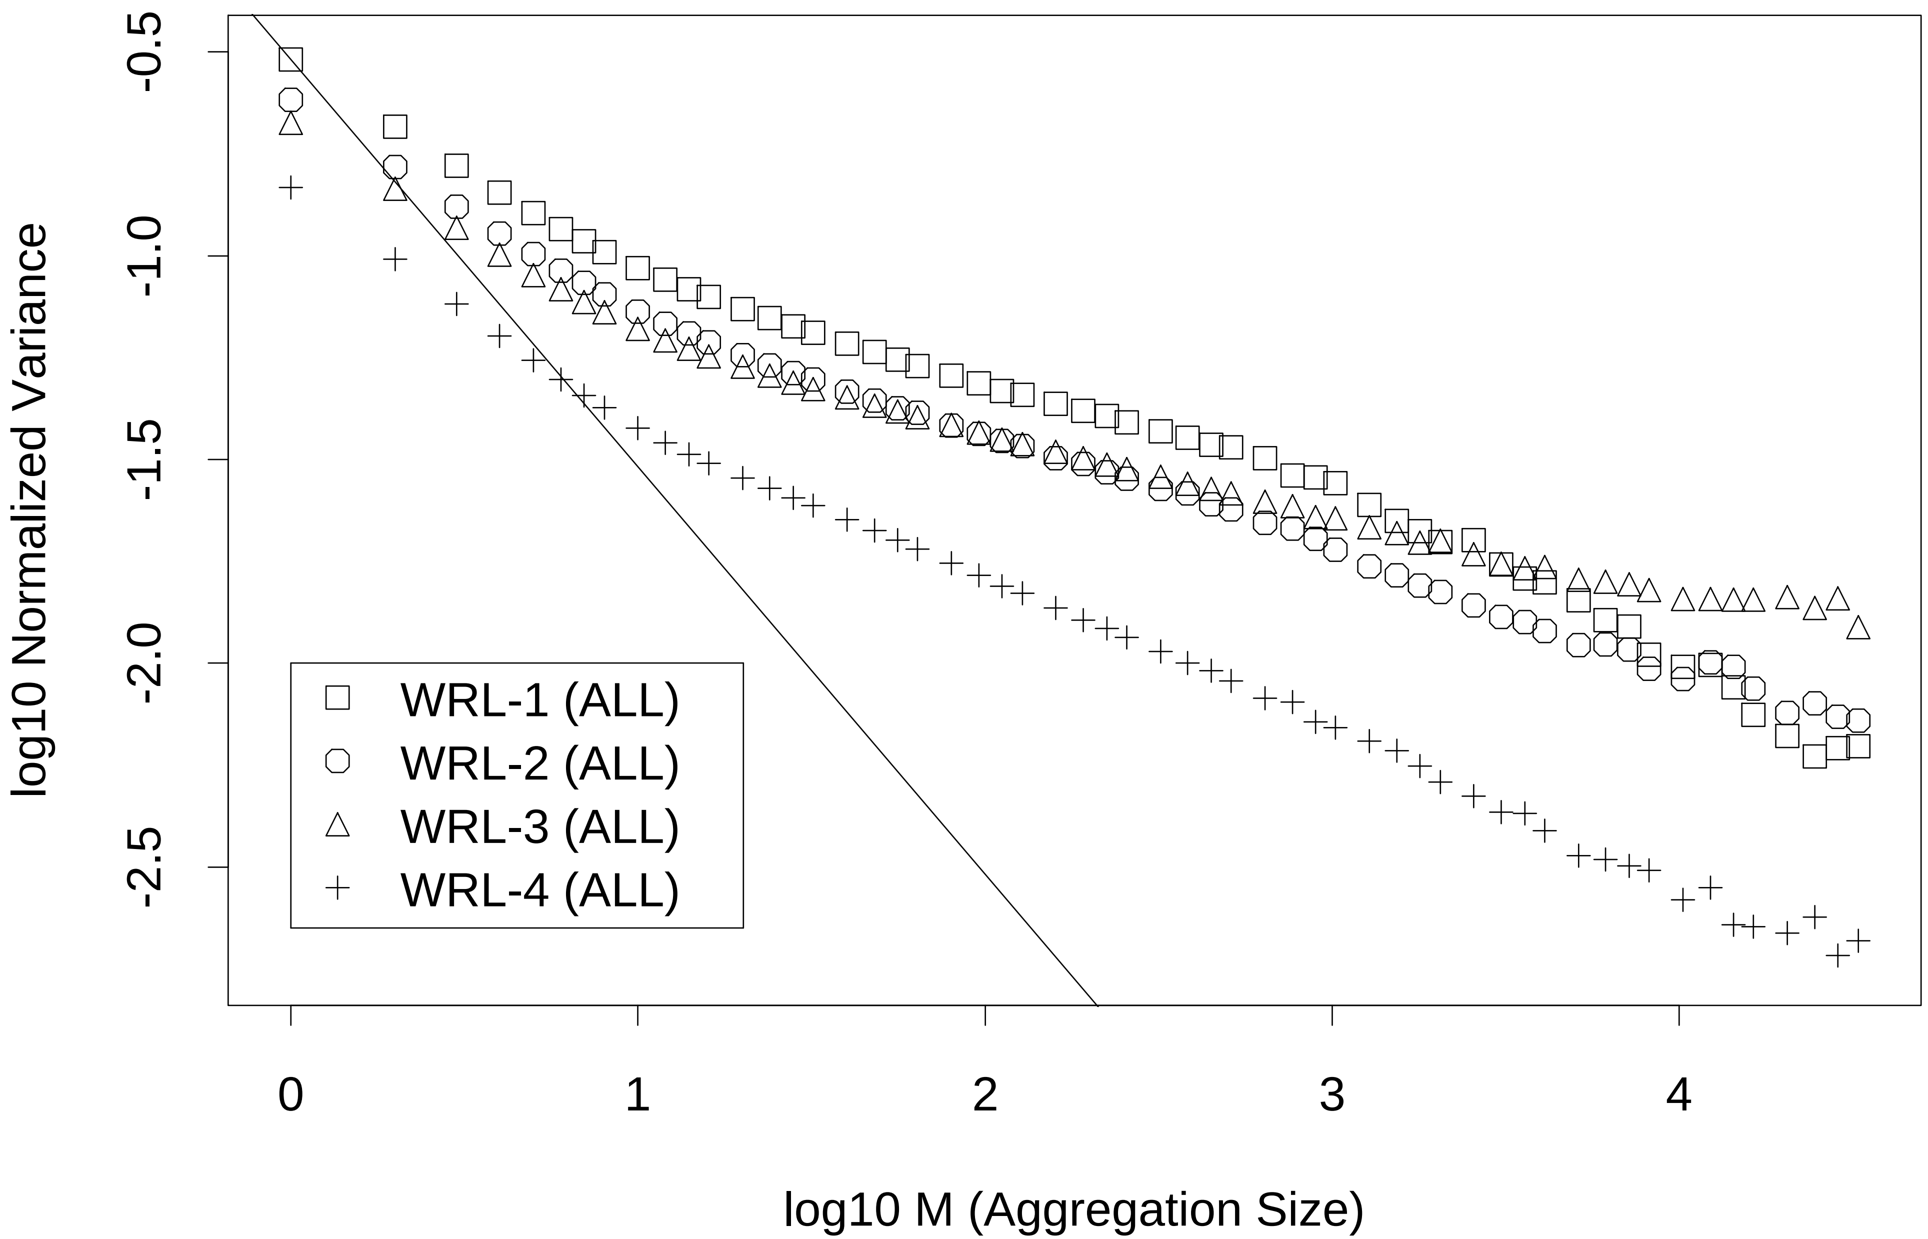

log10 Normalized Variance
-0.5
-1.0
-1.5
-2.0
-2.5
0
1
2
3
4
log10 M (Aggregation Size)
WRL-1 (ALL)
WRL-2 (ALL)
WRL-3 (ALL)
WRL-4 (ALL)

A look at fractals:

Informally: a (large) class of objects that have surprising scaling properties.

Classic example: How long is the British coastline? Answer: it depends on the length of your ruler.

The fractal scaling property:

There is no natural coastline “length.” If you look in different reference books, you get *different answers.*

Just as there is no natural network burst “length.”

Self-similar processes:

A special type of fractal is a *self-similar* process. Informally: something that "looks the same" regardless of scale. More formally: a stochastic process for which the autocorrelation of $X_k{}^{(}m)$ is the same as that for X_k.

Properties: "long-range dependence," or "long memory." Activity at all time scales.

Signature: variance decays slowly as scale increases.

$$\mathsf{Var}(X_k{}^{(m)}) \sim m^{-\beta}.$$

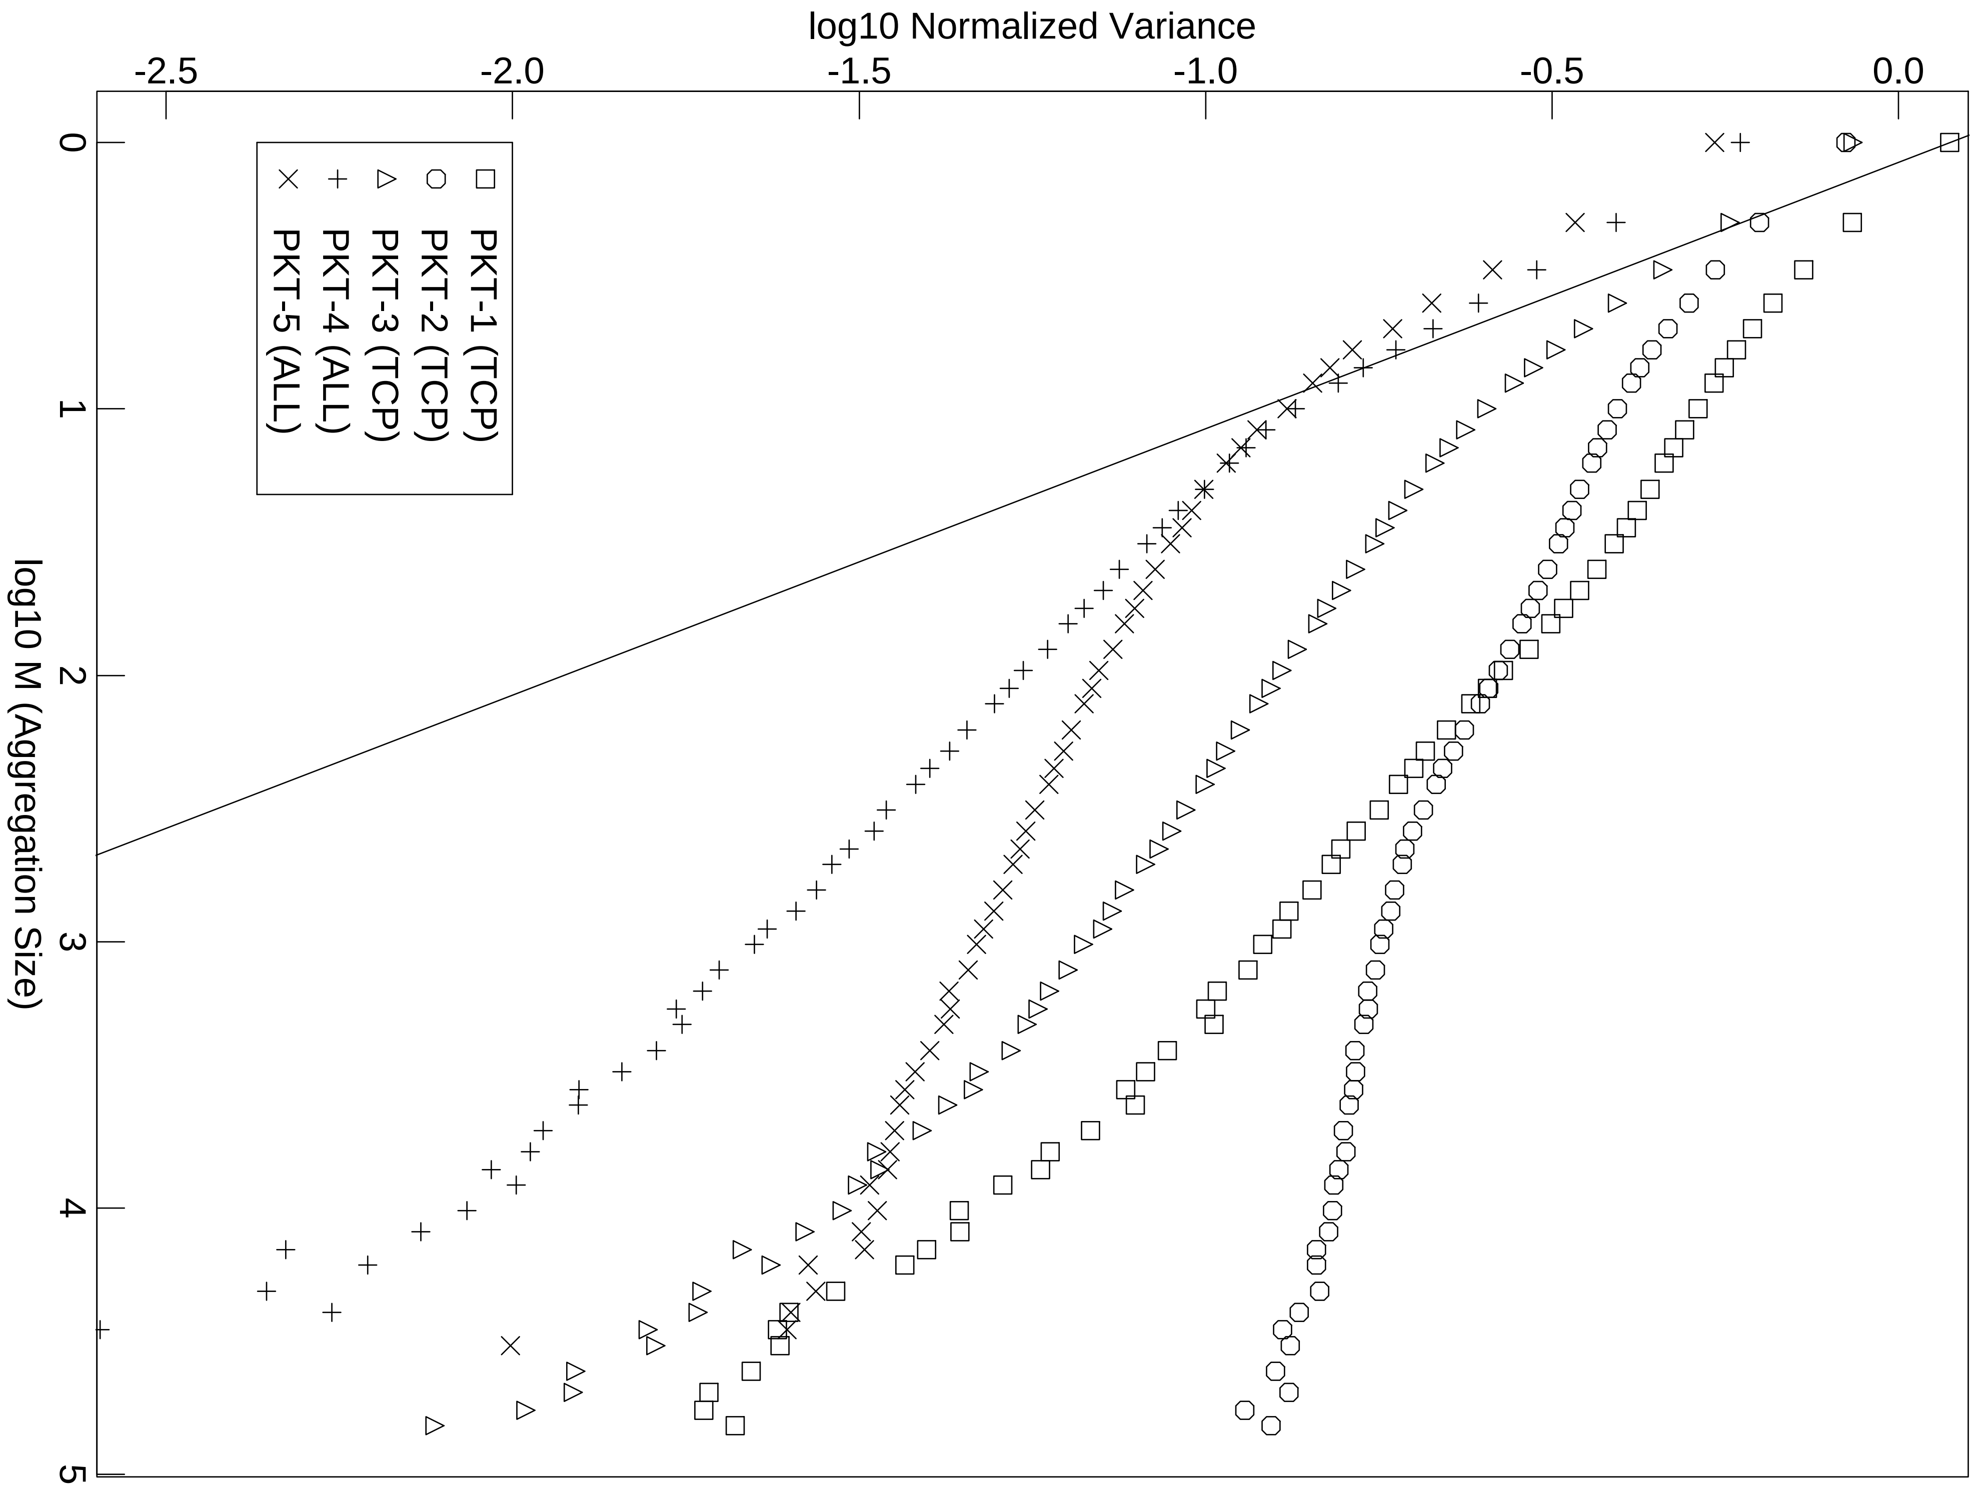
log10 Normalized Variance
-2.5
-2.0
-1.5
-1.0
-0.5
0.0
log10 M (Aggregation Size)
0
1
2
3
4
5
PKT-1 (TCP)
PKT-2 (TCP)
PKT-3 (TCP)
PKT-4 (ALL)
PKT-5 (ALL)

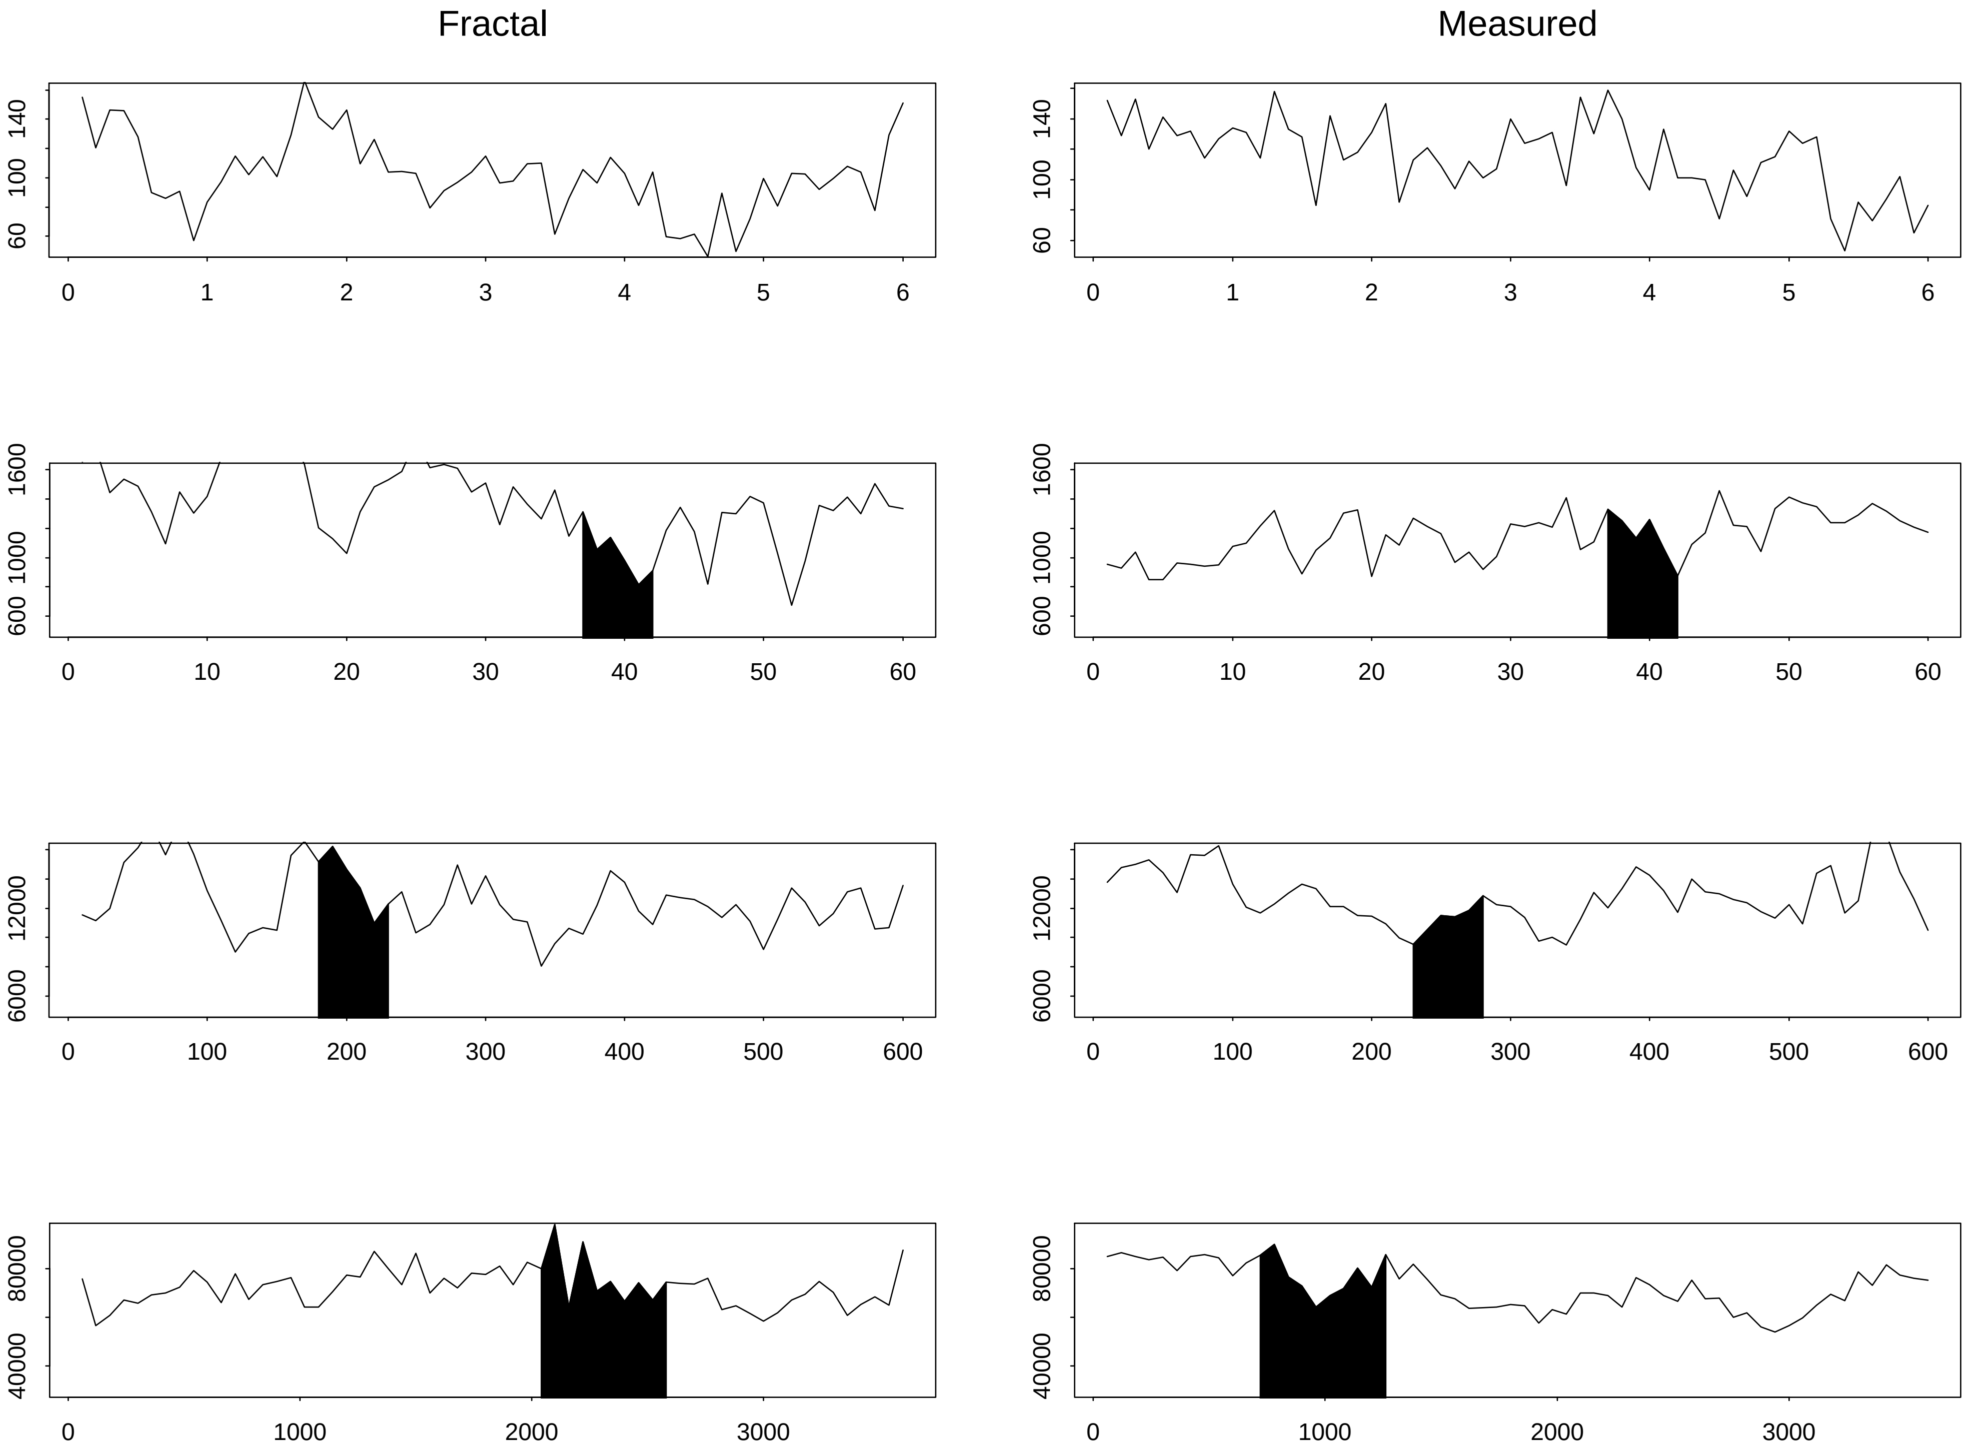

Fractal
Measured
60 100 140
0 1 2 3 4 5 6
600 1000 1600
0 10 20 30 40 50 60
6000 12000
0 100 200 300 400 500 600
40000 80000
0 1000 2000 3000

Recent network measurement & analysis:

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'.*

1993 Leland, Taqqu, Willinger & Wilson analyze additional LAN traffic. Propose self-similar models. Bombshell!

1993 Garrett studies coded Star Wars video: finds long-range dependence in corresponding bit-rate.

Recent network measurement & analysis, con't:

1993 Beran, Sherman, Taqqu & Willinger: video is replete with long-range dependence.

1994 Paxson & Floyd study WAN traffic. Find much closer to self-similar than any Poisson models. From *skeptic* $\Rightarrow$ *proponent.*

Others: LAN, NSFnet backbone, office automation network traffic all much closer matched by self-similar models.

How can this come about?

A key notion: *heavy* or *Pareto* tails:

$$F(x) = P[X \leq x] = 1 - (a/x)^{\beta}, \quad a, \beta \geq 0, \quad x \geq a.$$

For $\beta \leq 2$, a Pareto distribution has *infinite variance*. For $\beta \leq 1$, a Pareto distribution has *infinite mean*.

Implications of Pareto distributions:

The Pareto distribution is what you get if the *logarithm* of what you are measuring has an exponential distribution. That is, the *logarithm* is "memoryless" $\Rightarrow$ the thing you are measuring is independent of scale!

$\Rightarrow$ There are theorems directly relating processes having Pareto durations to self-similar behavior.

Signature: a straight line on a log-log distribution plot.

Heavy tails in computing:

CPU time consumed by Unix processes: [Leland & Ott, 1986]

Size of Unix files: [Irlam, 1993]

Size of compressed video frames: [Garrett, 1993]

Size of FTP bursts: [Paxson & Floyd, 1994]

Telnet packet interarrivals: [Danzig et al, 1991], [Paxson & Floyd, 1994]

Ethernet bursts: [Willinger et al, 1995]

Size of Web items: [Crovella & Bestavros, 1996]

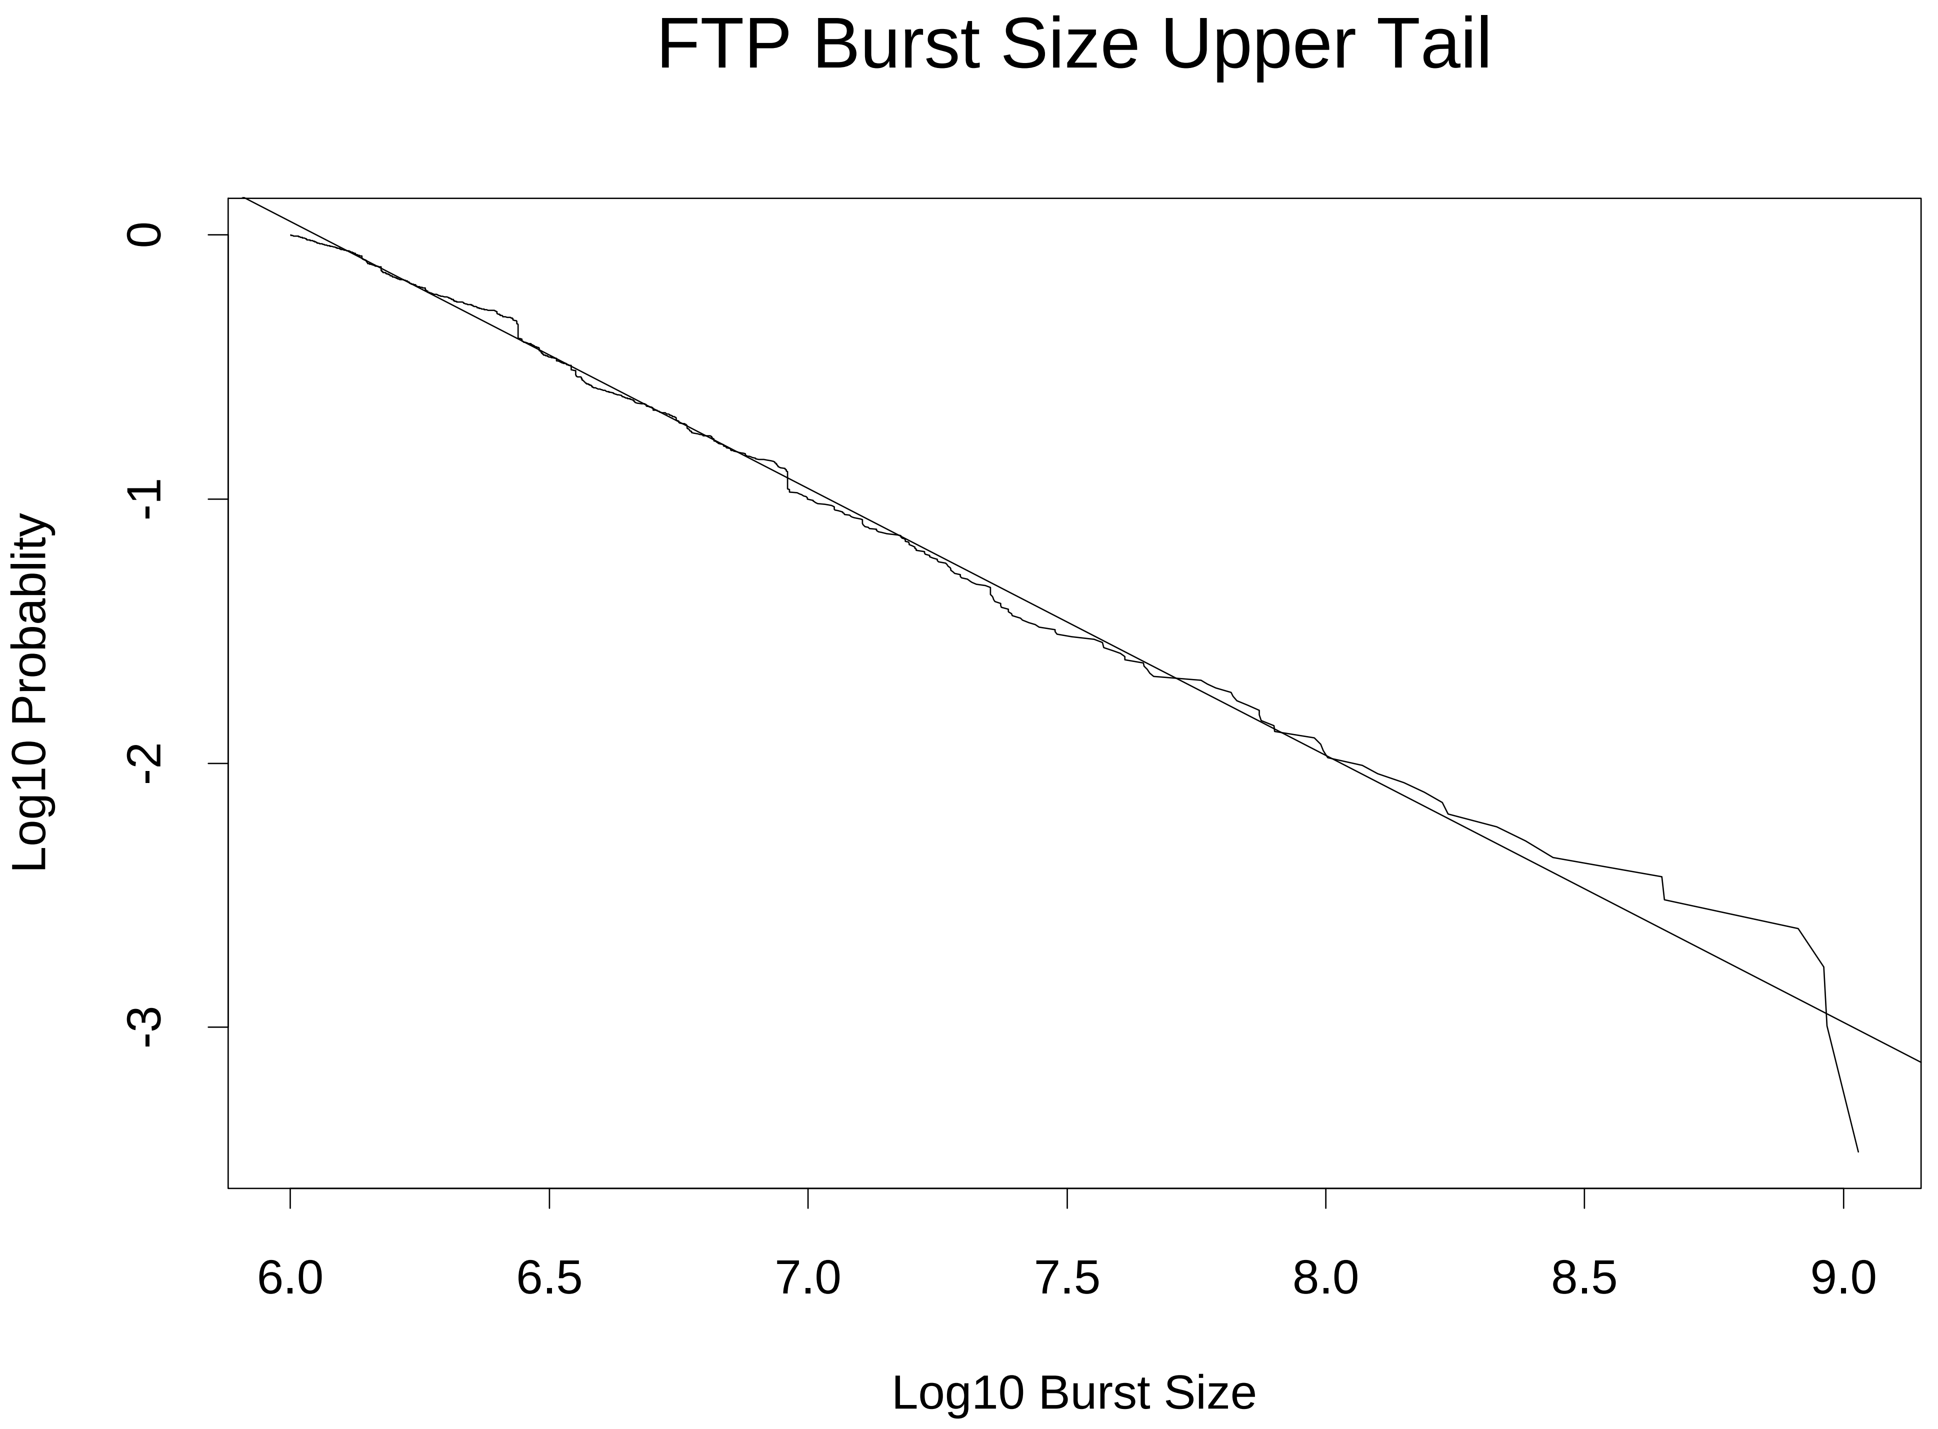
FTP Burst Size Upper Tail
0
-1
-2
-3
Log10 Probablity
6.0
6.5
7.0
7.5
8.0
8.5
9.0
Log10 Burst Size

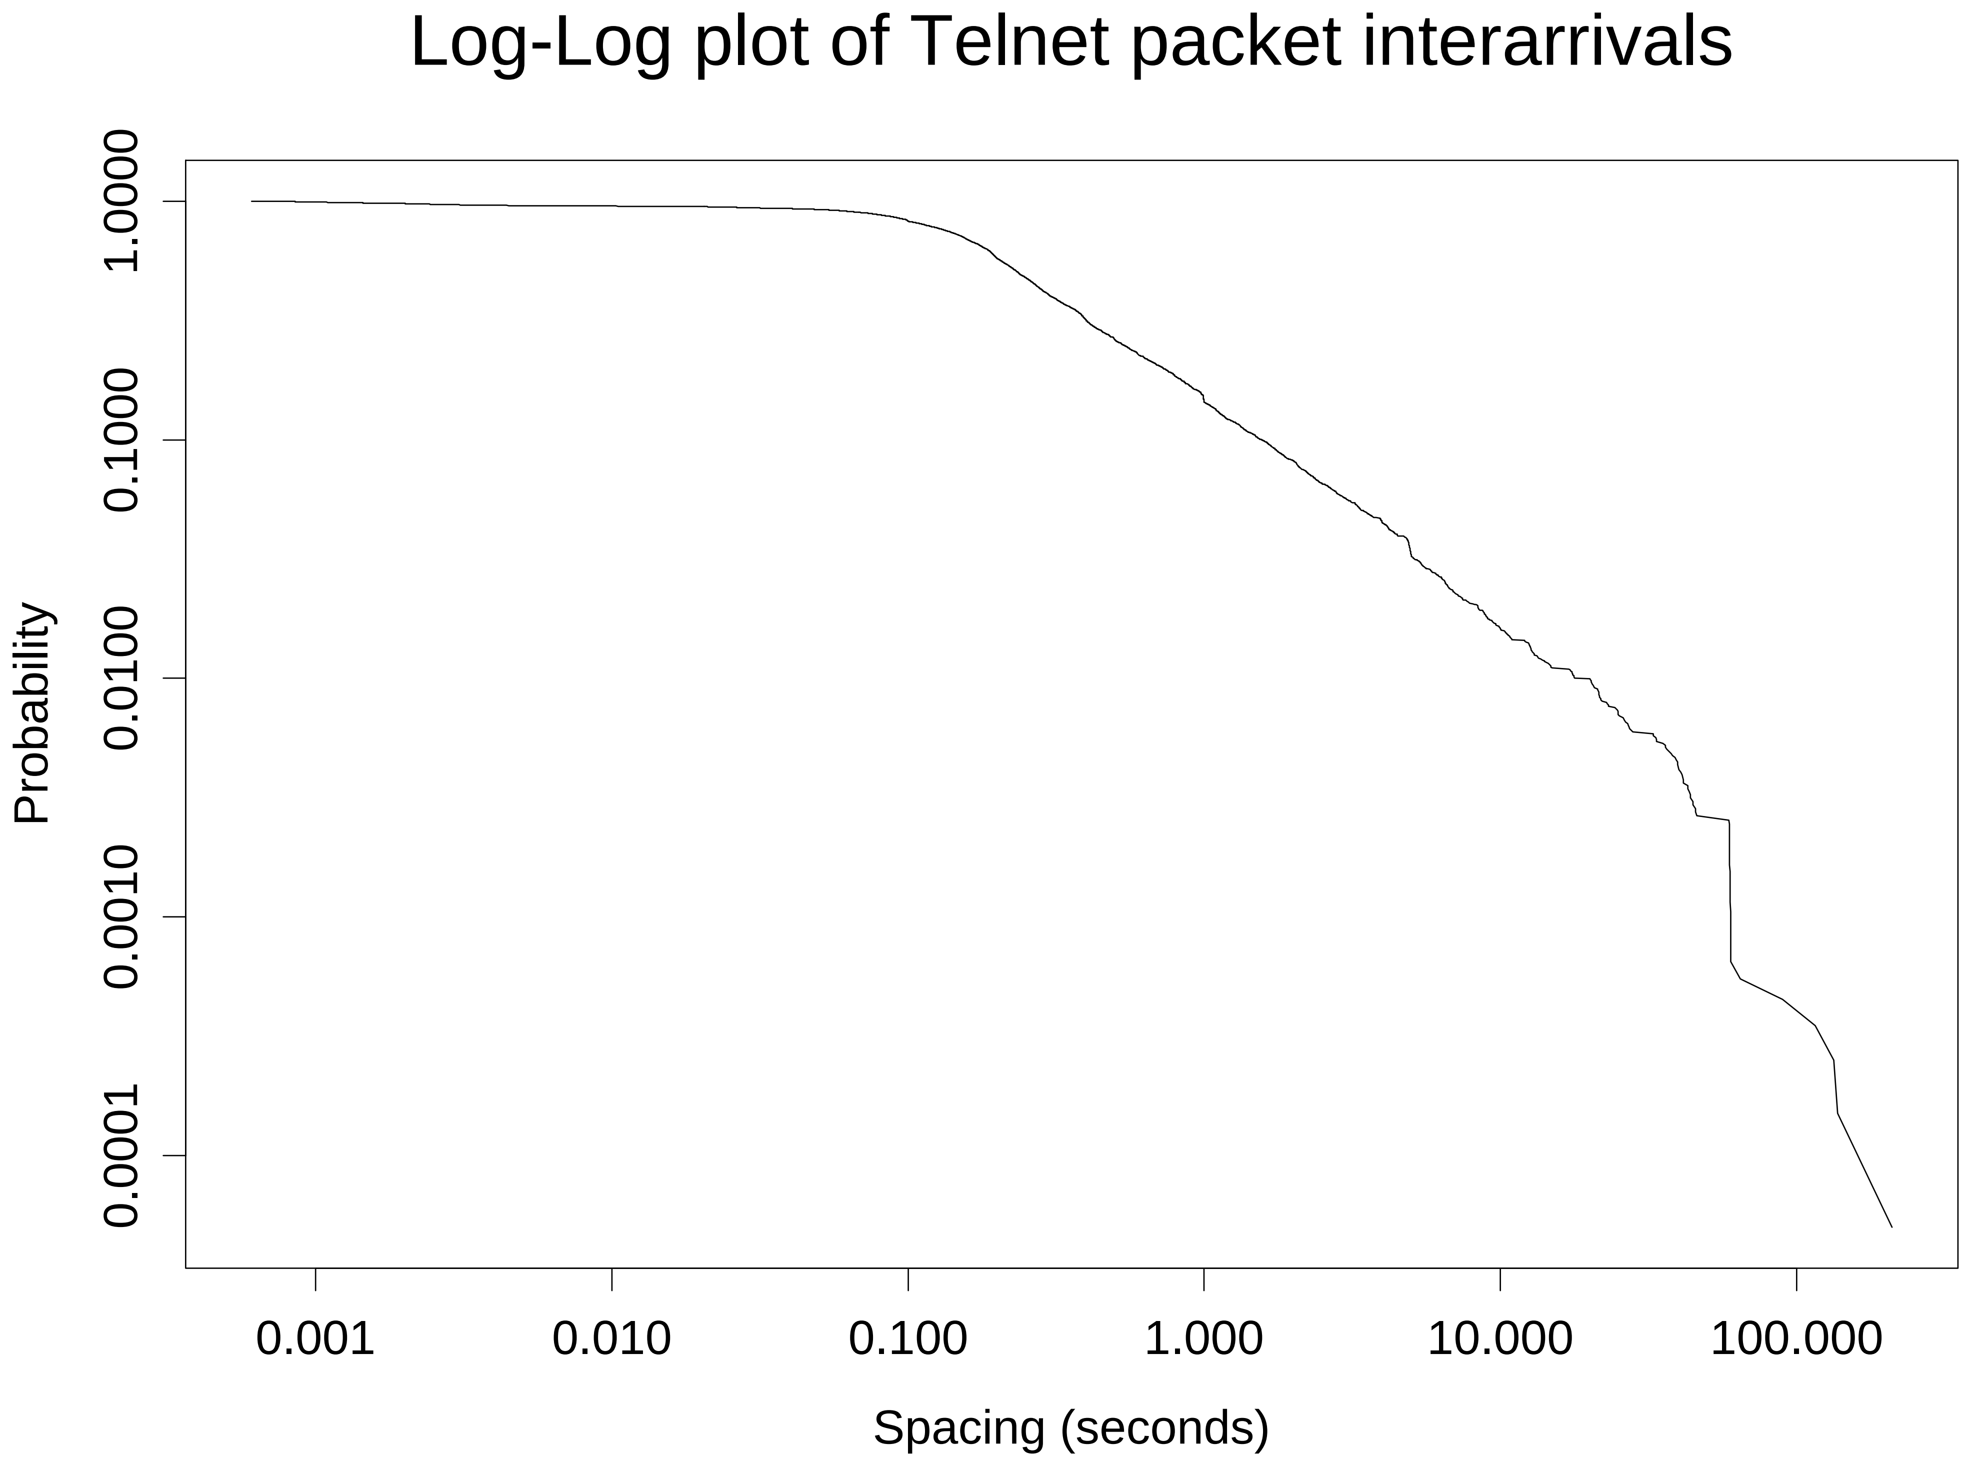
Log-Log plot of Telnet packet interarrivals
Probability
1.0000
0.1000
0.0100
0.0010
0.0001
0.001
0.010
0.100
1.000
10.000
100.000
Spacing (seconds)

What does this mean for network traffic?

Queueing delays are *much higher* in the presence of long-range dependence than for Poisson traffic (Erramilli, Narayan, Willinger).

To avoid dropping packets, buffers have to be huge.

You have to be very careful predicting future traffic based past measurements.

You can't look at a little bit of video and decide how much buffer it's going to require.

Major research questions for fractal modeling:

How do we incorporate short-range dependence? What is the right model for short-range dependence, anyway?

How do we model <u>interarrivals</u> given that we know <u>arrival counts</u>?

<u>Lots</u> of remaining work in understanding causes, implications, when it matters & when it doesn't, good measurement techniques, ways to ameliorate, etc.

What we've learned, where we are:

“[Hurst put us] in one of those situations, so salutary for theoreticians, in which empirical discoveries stubbornly refuse to accord with theory.” - E.H. Lloyd

$\Rightarrow$ Measure, measure, measure . . .

. . . and *analyze*.

Expect heavy tails — think about how to accommodate them in your designs.

And as soon as you find heavy tails, think *fractal*.

Bibliography:

J. Beran, R. Sherman, M. Taqqu, and W. Willinger, “Variable-Bit-Rate Video Traffic and Long-Range Dependence,” *IEEE Transactions on Communications*, 43(2-4), pp. 1566-1579, 1995.

M. Crovella and A. Bestavros, “Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes,” to appear in *Proceedings of SIGMETRICS '96.*

R. B. D'Agostino and M. A. Stephens, editors, “Goodness-of-Fit Techniques”, Marcel Dekker, Inc., 1986.

P. Danzig and S. Jamin, “tcplib: A Library of TCP Internetwork Traffic Characteristics”, *Report CS-SYS-91-01*, Computer Science Department, University of Southern California, 1991.

Bibliography, con't:

P. Danzig, S. Jamin, R. Cáceres, D. Mitzel, and D. Estrin, “An Empirical Workload Model for Driving Wide-area TCP/IP Network Simulations”, *Internetworking: Research and Experience*, 3 (1), pp. 1-26, 1992.

H. Fowler and W. Leland, “Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management,” *IEEE JSAC*, 9(7), pp. 1139-1149, September, 1991.

M. Garrett and W. Willinger, “Analysis, Modeling and Generation of Self-Similar VBR Video Traffic,” *Proceedings of SIGCOMM '94*, pp. 269-280, September, 1994.

R. Gusella, “A Measurement Study of Diskless Workstation Traffic on an Ethernet,” *IEEE Transactions on Communications*, 38(9), pp. 1557-1568, September, 1990.

Bibliography, con't:

G. Irlam (gordoni@netcom.com), “ufs'93 [Updated file size survey results],” USENET newsgroup comp.os.re-search, message 2ddp3b$jn5@darkstar.UCSC.EDU, Nov. 29, 1993.

V. Jacobson, C. Leres, and S. McCanne, *tcpdump*, available via anonymous ftp to ftp.ee.lbl.gov, June, 1989.

R. Jain and S. Routhier, “Packet Trains — Measurements and a New Model for Computer Network Traffic,” *IEEE JSAC*, 4(6), pp. 986-995, September, 1986.

W. Leland and T. Ott, “Load-balancing Heuristics and Process Behavior,” *PERFORMANCE '86 and ACM SIGMETRICS 1986 Joint conference on Computer Performance Modelling, Measurement and Evaluation*, North Carolina State University, pp. 54-69, May 1986.

Bibliography, con't:

W. Leland, M. Taqqu, W. Willinger, and D. Wilson, “On the Self-Similar Nature of Ethernet Traffic”, Proceedings of SIGCOMM '93, September, 1993, pp. 183-193.

W. Leland, M. Taqqu, W. Willinger, and D. Wilson, “On the Self-Similar Nature of Ethernet Traffic (Extended Version)”, *IEEE/ACM Transactions on Networking*, 2(1), pp. 1-15, February 1994.

P. Martin-Löf, “The Notion of Redundancy and Its Use as a Quantitative Measure of the Discrepancy between a Statistical Hypothesis and a Set of Observational Data”, *Scandinavian Journal of Statistics*, 1(1), pp. 3-18, 1974.

Bibliography, con't:

P. Pawlita, “Two Decades of Data Traffic Measurements: A Survey of Published Results, Experiences and Applicability”, Teletraffic Science for New Cost-Effective Systems, Networks and Services, ITC-12, M. Bonatti (Editor), Elsevier Science Publishers B.V. (North-Holland), 1989.

V. Paxson, “Empirically-Derived Analytic Models of Wide-Area TCP Connections,” *IEEE/ACM Transactions on Networking*, 2(4), pp. 316-336, August, 1994; `ftp://ftp.ee.lbl.gov/papers/WAN-TCP-models.ps.Z`.

V. Paxson, “Growth Trends in Wide-Area TCP Connections”, *IEEE Network*, 8(4), pp. 8-17, July/August 1994; `ftp://ftp.ee.lbl.gov/papers/WAN-TCP-growth-trends.ps.Z`.

Bibliography, con't:

V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", *Proceedings of SIGCOMM '94*, pp. 257-268, September 1994. A revised version appears in *IEEE/ACM Transactions on Networking*, 3(3), pp. 226-244, June, 1995; `ftp://ftp.ee.lbl.gov/papers/WAN-poisson.ps.Z`.

S. Pederson and M. Johnson, "Estimating Model Discrepancy", *Technometrics*, 32(3), August, 1990, pp. 305-314.

W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, "Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *Proceedings of SIGCOMM '95*, pp. 100-113, September 1995.