Statistically Accurate Network Measurements

Matthew Roughan
<matthew.roughan@adelaide.edu.au>

with Hung Nguyen
Applied Mathematics
School of Mathematical Sciences
University of Adelaide

June 4, 2010

Network Measurements

"Measure what is measurable, and make measurable what is not so."

Galileo Galilei, (1564 - 1642)

"To measure is to know."

Lord Kelvin, Sir William Thomson (1824-1907)

The problem

- Computer science
  - we deal with digital data
  - there is a tendency to believe we see reality
- Measurements are **ALWAYS** flawed
  - measurement precision
  - ambiguity
  - missing data
  - measurement artifacts
  - impact of the measurements themselves
- So what do I do?
  - make more measurements
  - do a bit of stats ...

So What?

Statistics has been solving these sorts of problems for a while.

- Why isn't this all easy?

So What?

Statistics has been solving these sorts of problems for a while.

- Why isn't this all easy?
- I'm a bit clueless

So What?

Statistics has been solving these sorts of problems for a while.

- Why isn't this all easy?
- I'm a bit clueless
- There are some hard bits
  - non-Gaussian distributions and heavy-tails
  - strong correlations (and long-range dependence)
  - structural biases in sampling
  - we don't get to see what we would like to see

Examples: Traffic

- estimating traffic matrices from link data [1, 2, 3, 4]
- interpolating missing data in traffic matrices [5]
- SNMP link measurement errors [6]

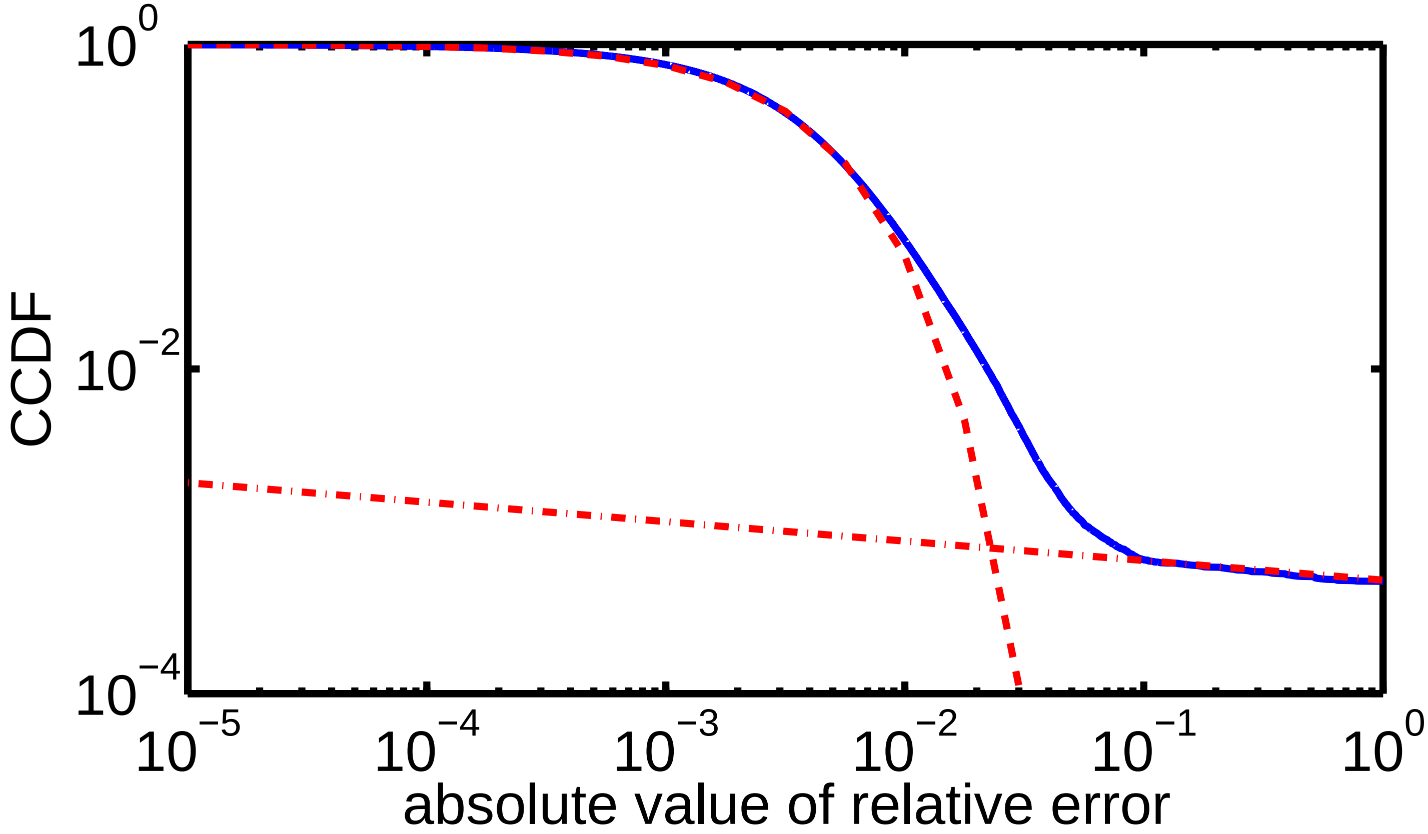

Examples: Routing and Topology

- traceroute aliasing problem [7]
- load-balancing [8]
- bias in sampled viewpoint [9, 10, 11]

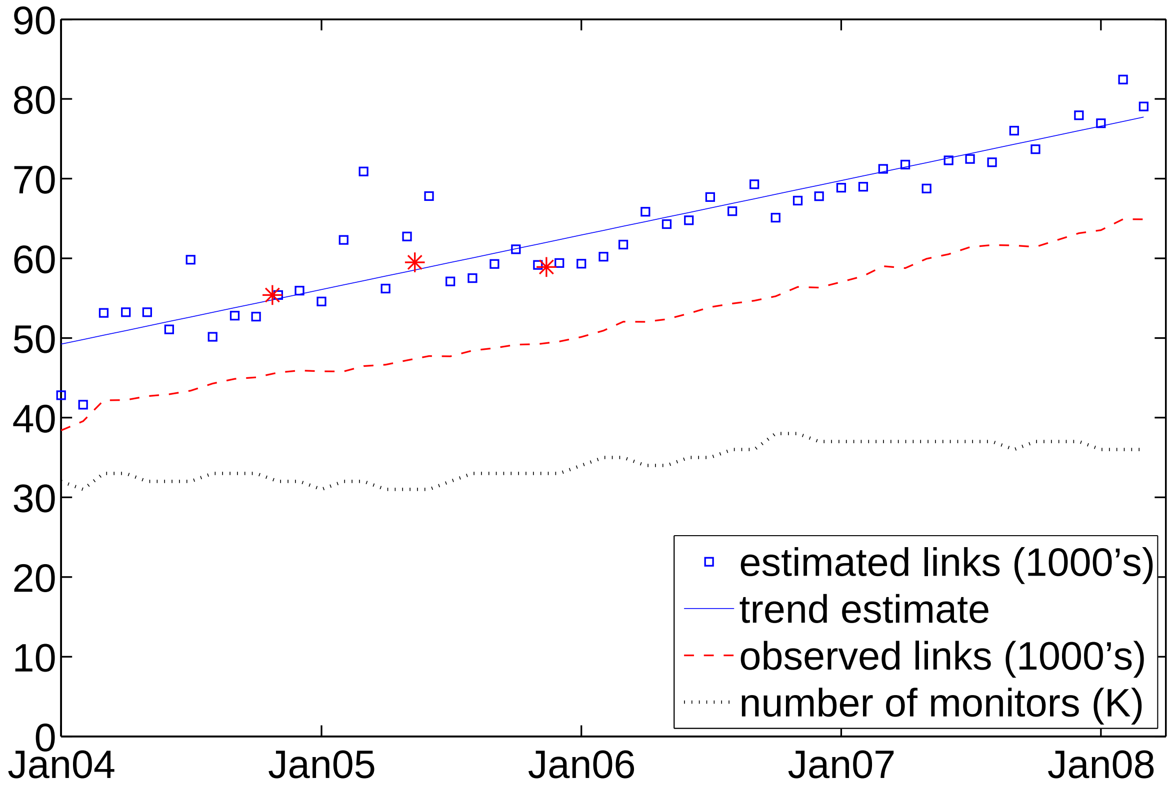

Examples: Performance

- Active performance measurements
- Send probe packets from $A \rightarrow B$ across the network
- Measure the performance experienced by packets

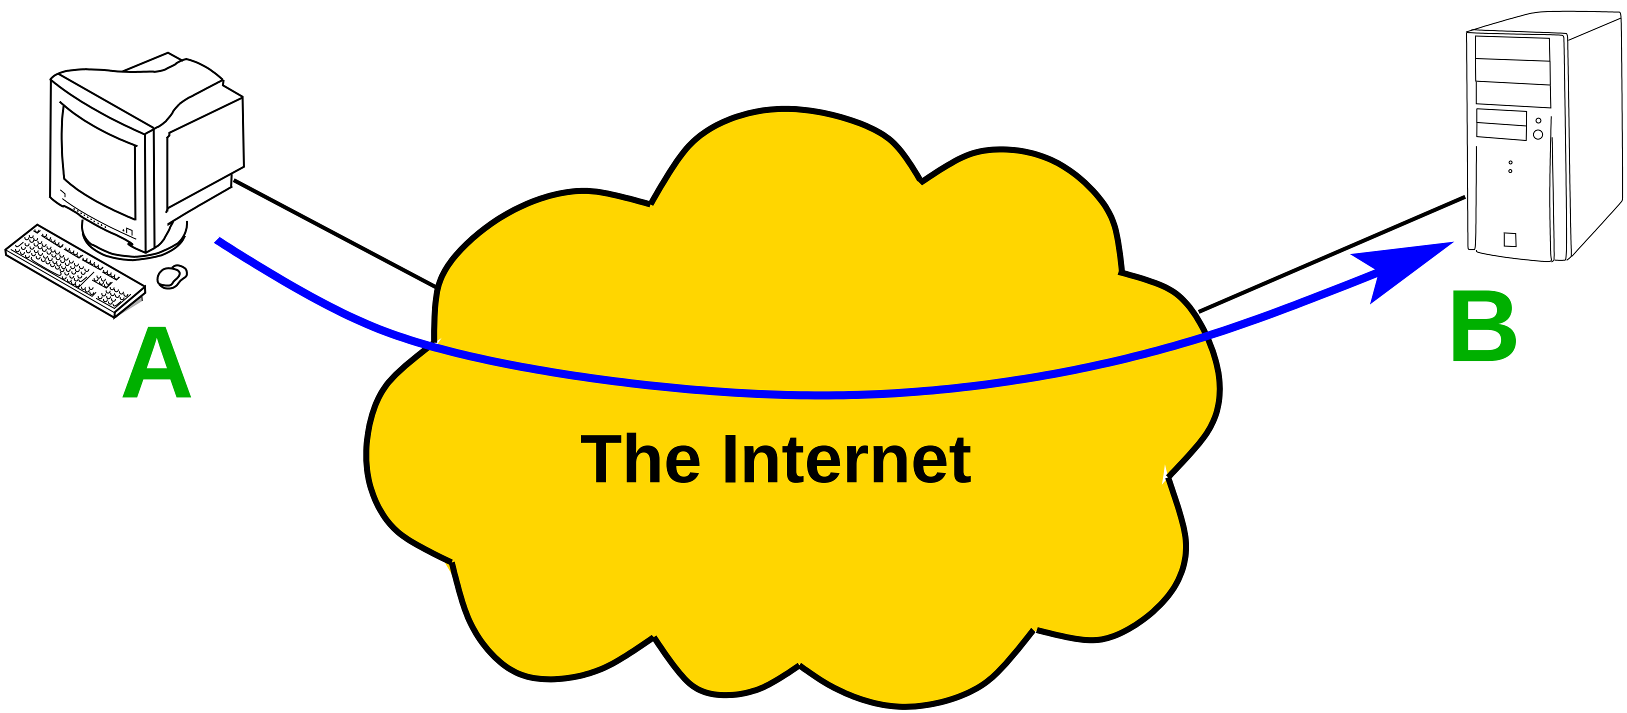

The problem

- How accurate are the probes?
- How many probes should we send?
- When is there a network problem?

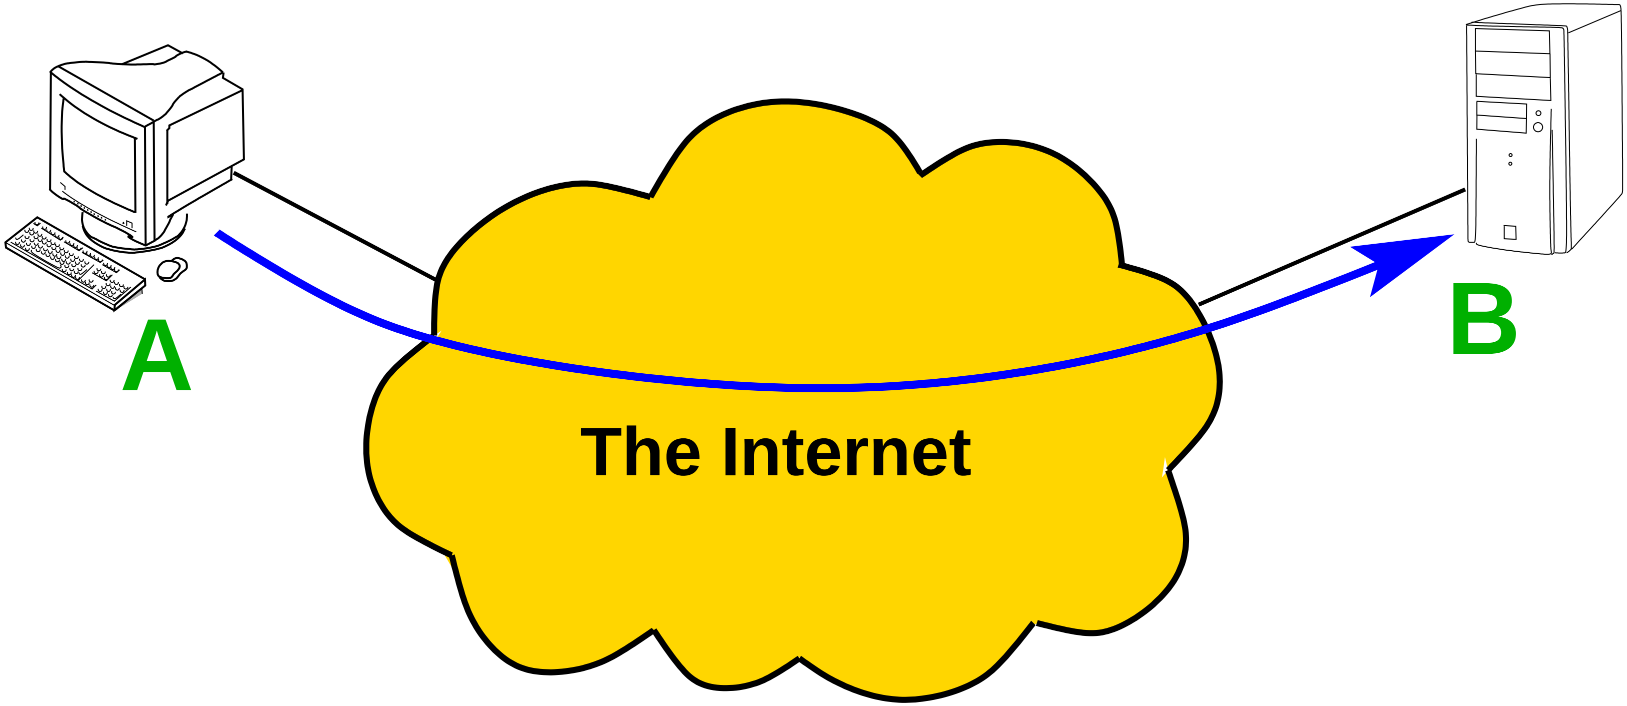

Accuracy

What do we mean by accuracy?

- probes are samples from some underlying process
  - assume measurement equipment has no flaws
- measurements over time can be averaged to give estimates
- implicit assumption of stationary ergodic process $\Rightarrow$ a time average converges to an ensemble average
- variance can be estimated
  - typically under the assumption of independent measurements

Accuracy of **estimates** not individual measurements

The plural of ...

"The plural of anecdote is not evidence."
Alan L. Leshner

The plural of ...

"The plural of anecdote is not evidence."
Alan L. Leshner

A Unicorn

If you Google "Unicorn", you'll get about 3 million images, mostly resembling this.

The plural of ...

"The plural of anecdote is not evidence."
Alan L. Leshner

A Rhino

Unicorn images aren't too evocative of the Rhinoceros

The plural of ...

"The plural of anecdote is not evidence."
Alan L. Leshner

A NarWhal

Throw in some narwhal horns for confusion

Correlated Measurements

The underlying problem is that correlated measurements don't give you extra information in the same way independent measurements do.

- if the measurements are really from the same source (e.g. Unicorn myths) then they give you no extra information
- what about packet probes?
  - are they correlated?

Answers

- Yes!
  - there are correlations in measurements, e.g.,[12, 13, 14, 15]
  - the short-range correlations can be quite strong even where there is no long-range dependence.
  - these correlations must be taken into account in estimating accuracy, confidence intervals, ...
- correlations get stronger the closer measurements are in time [16]
  - within a finite sampling period there is no point in arbitrarily increasing the number of samples
- correlations are load dependent [16]

Consider loss measurements

- Used for
  - alarms – signal network problems
  - Service-Level Agreement (SLA) compliance
- typical approach:
  - assume Bernoulli (independent) trials
  - but its well known there are correlations
- Why do we care about accuracy
  - are we compliant with an SLA
  - if "accuracy" of measurements is wrong, we may mistakenly claim compliance, or otherwise

Loss measurement estimation

Standard estimator of loss probability

$$\hat{p} = \frac{1}{N} \sum_{i=1}^{N} I_i, \qquad (0)$$

where $I_i = 1$ if probe ith is lost, and $I_i = 0$ otherwise, and

$$\text{Var}\left[\hat{p}\right] = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} R(\tau_{ij}),$$

- $R(\tau)$ is the autocovariance of the loss process.
- τ_{ij} is the time between the i and jth probes.

Possible Approaches

- Assume independence

$$\text{Var}\left[\hat{p}\right] = p(1-p)/N,$$

- Estimate R_{ij}:
  - treat measurements like a discrete process
  - but then we are limited to those measurements
  - better to estimate underlying process
- Estimate $R(\cdot)$ directly from data
  - estimates of $R(\cdot)$ have a high variance themselves
- **Estimate $R(\cdot)$ with the help of a model**

Our Approach

- Model + parametric estimation
  - model simple enough to work with
  - general enough to fit data
  - model underlying loss, not packet losses
- Markov model doesn't work too well
  - need too many states
- Alternating renewal model is attractive
  - approximately matches data
  - fits our intuition about loss bursts
    - from queueing
    - from routing events

Alternating Renewal Model

- On/Off process
  - On and Off times are independent
  - On times form IID sequence
  - Off times form IID sequence
- Parameterise the distribution of On/Off times
  - we used Gamma distribution
  - natural generalization of the exponential

$$\text{Prob}\{X \in [x, x+dx)\} = x^{k-1} \frac{e^{-x/\theta}}{\Gamma(k)\theta^k} dx,$$

- two parameters
  - shape k ($k = 1$ gives exponential distribution)
  - scale θ

SAIL

Samples from underlying loss-process

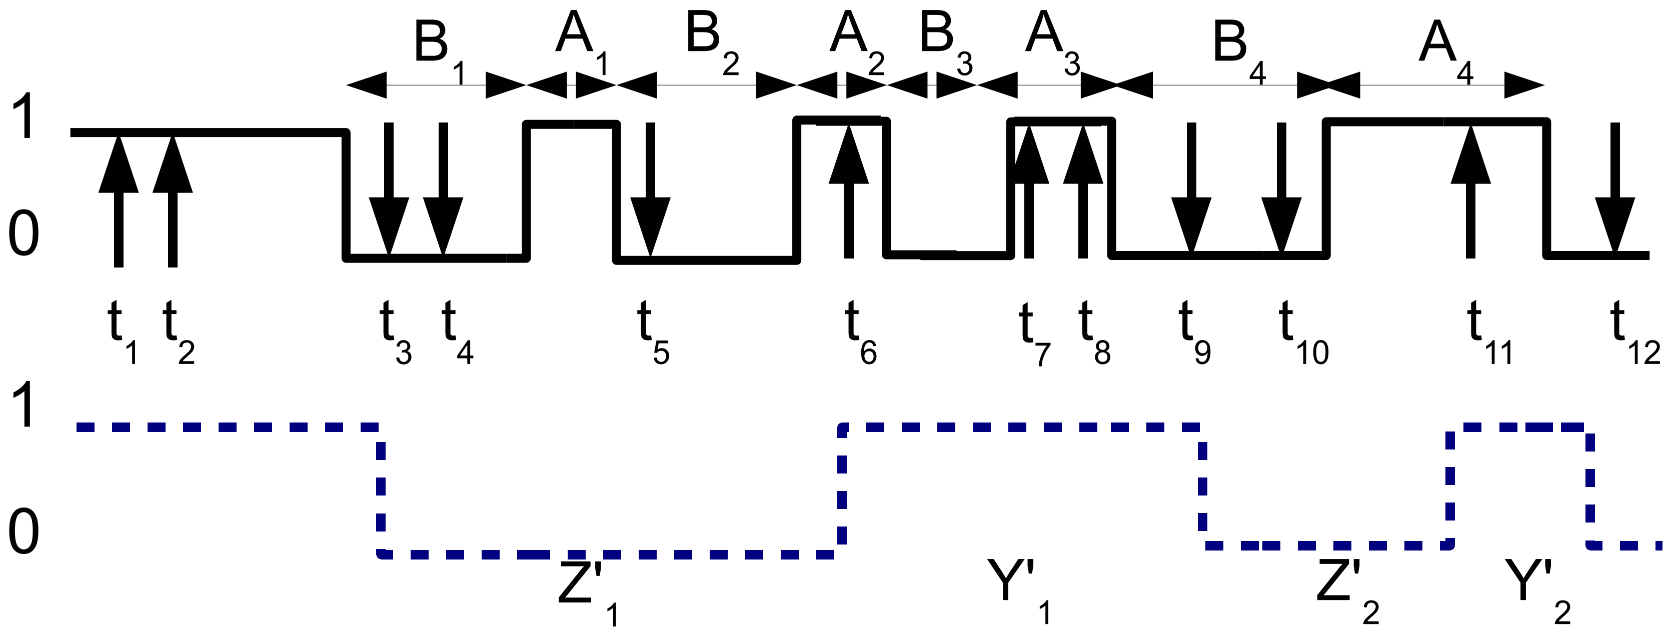

- samples can miss some intervals
- direct estimation of intervals overestimates them
- estimation using Hidden Semi-Markov Model techniques

Simulations: test estimates

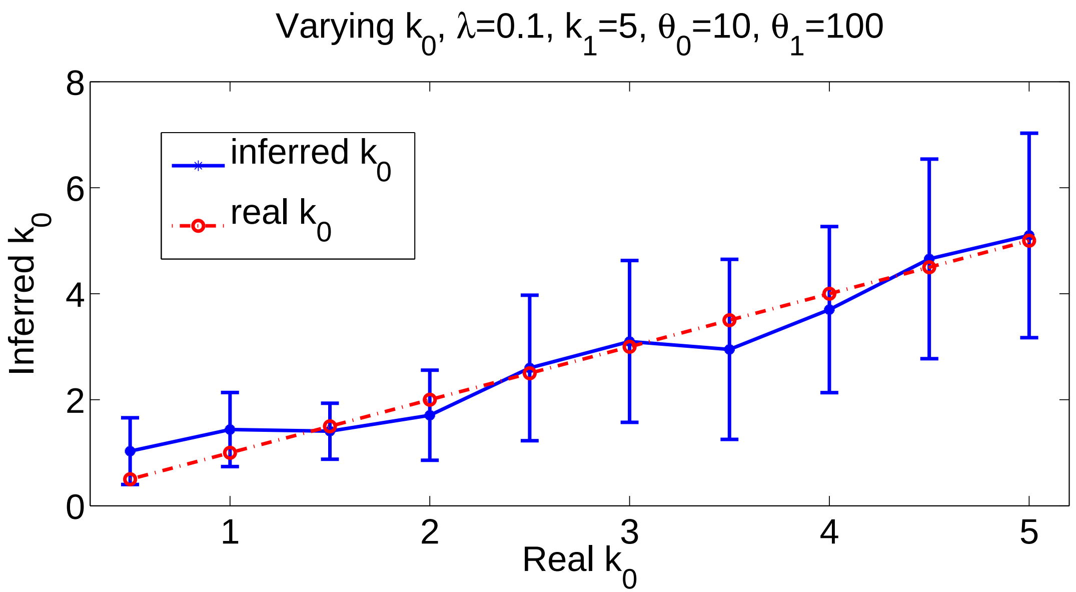

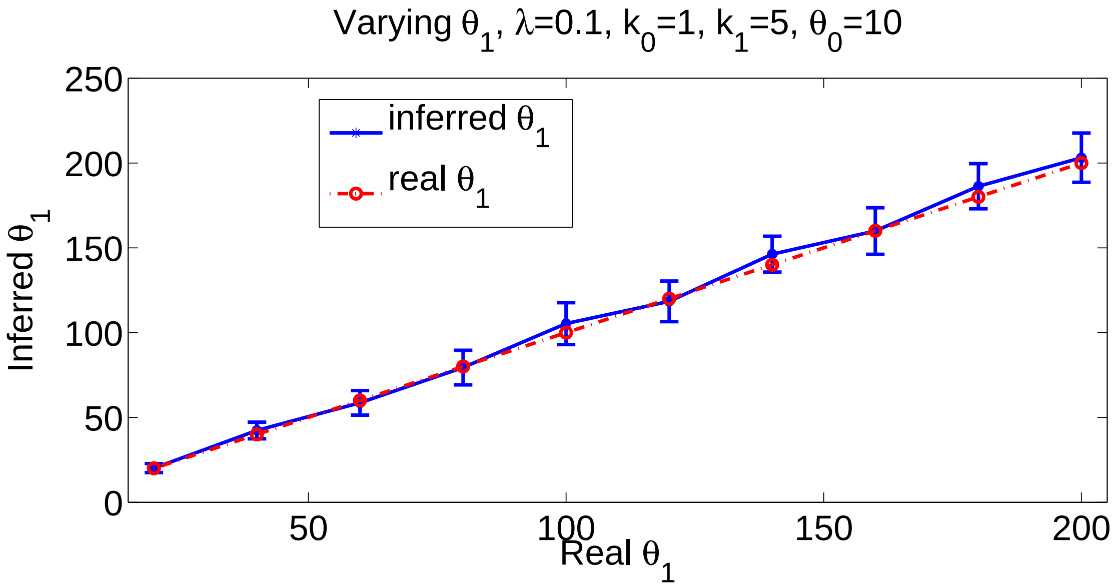

Experiments

- Measurement locations
  - One way Planetlab → Planetlab
  - Between dedicated machines
    - EPFL and Australia
  - Web server probes
  - DNS server probes (provided by QUEEN)
- Standard Traces
  - 10,000 40 byte UDP probe packets
  - Sent as Poisson process average rate 10 packets per second
  - Filtered for stationarity and non-negligible loss
  - 1,100 traces analysed

Results: θ

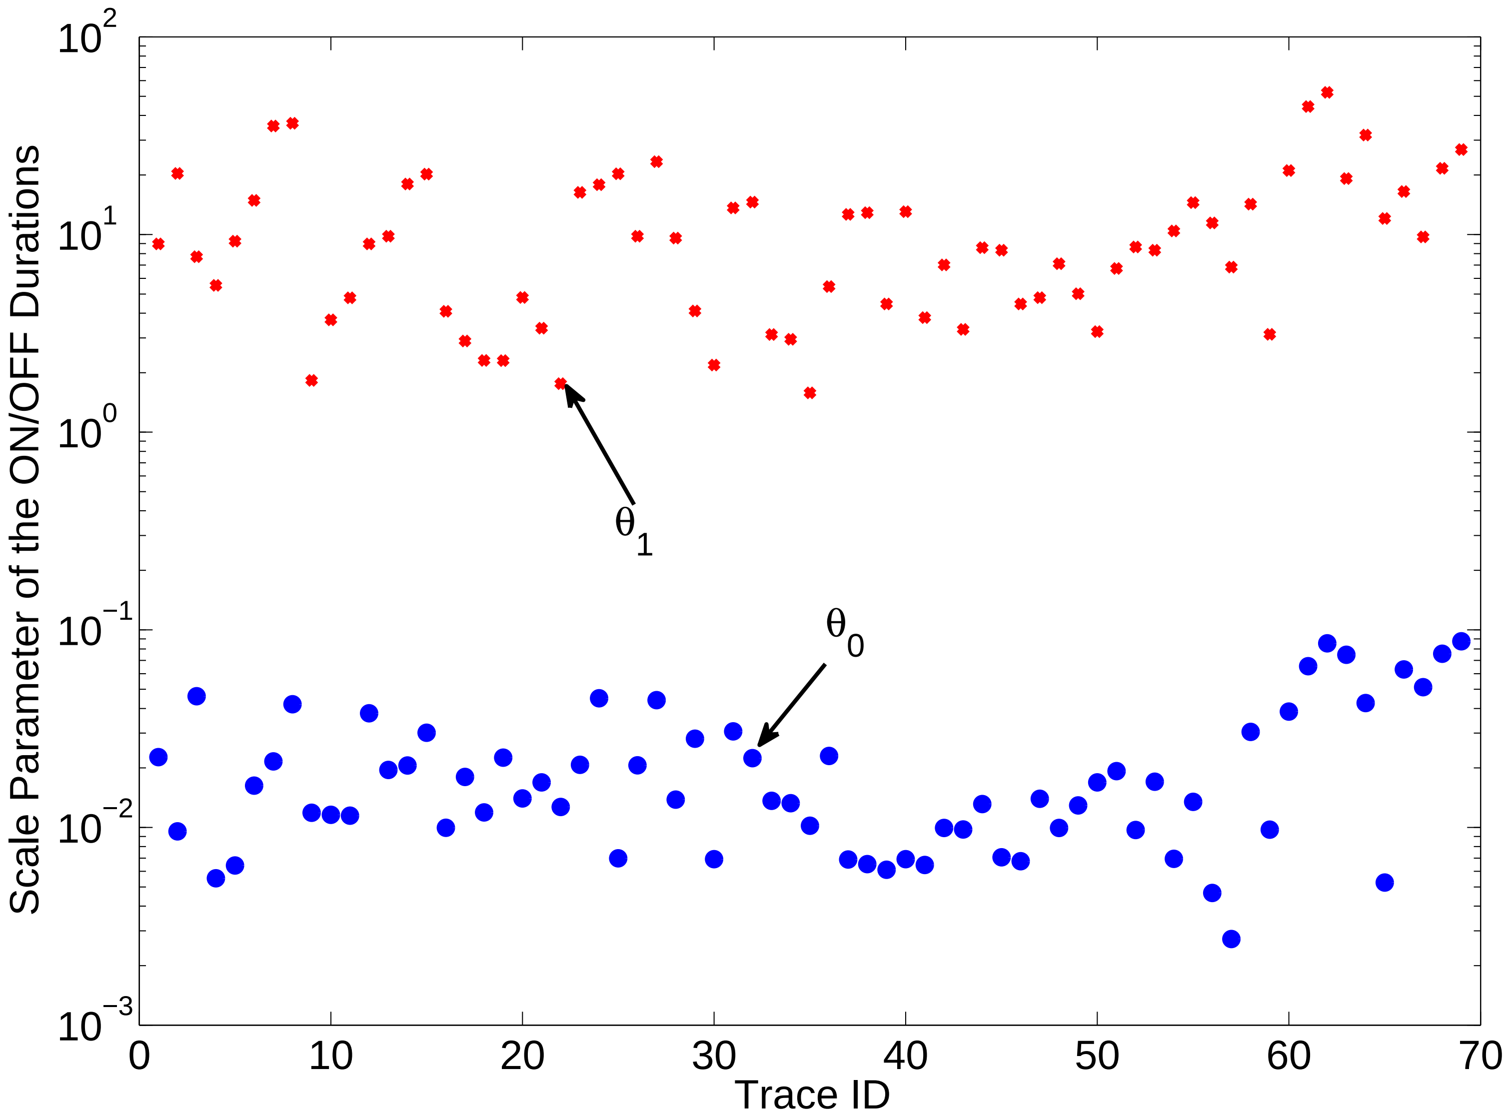

10^2
10^1
10^0
10^{-1}
10^{-2}
10^{-3}
Scale Parameter of the ON/OFF Durations
θ_1
θ_0
0
10
20
30
40
50
60
70
Trace ID

Results: k

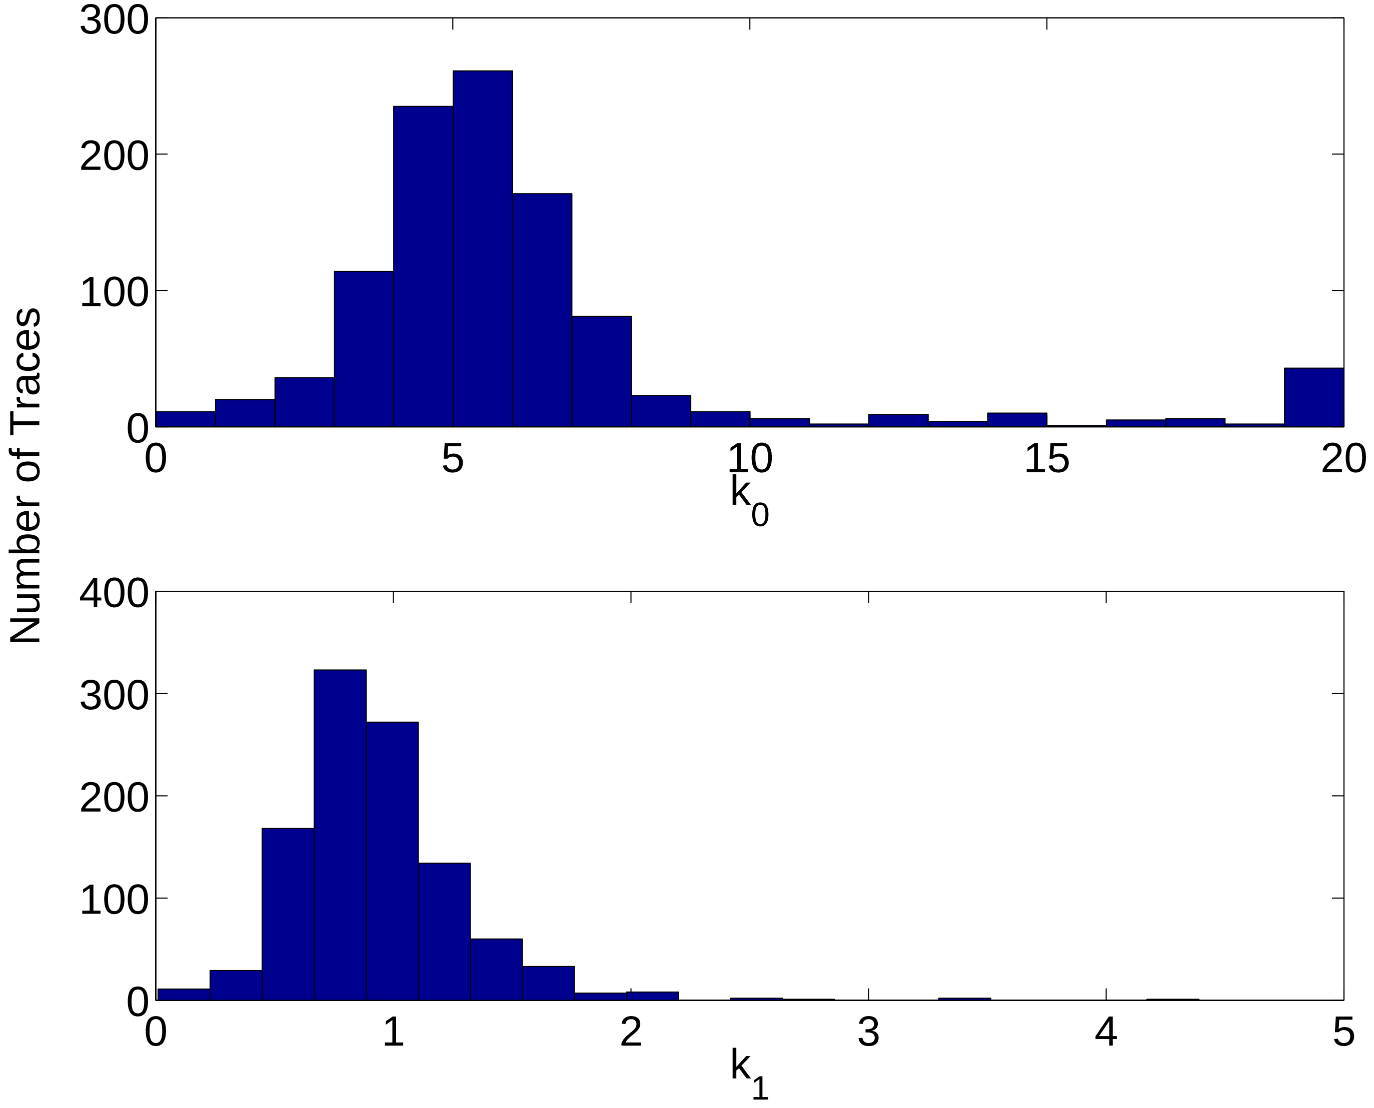

300
200
100
0
0
5
10
15
20
k0
Number of Traces
400
300
200
100
0
0
1
2
3
4
5
k1

Validation

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- Hard to get ground truth to compare against
- test by sub-sampling
  - randomly segregate data into two sections
    - each is still Poisson probe process
  - compute results for both
  - look for consistency

Validation: consistency variance

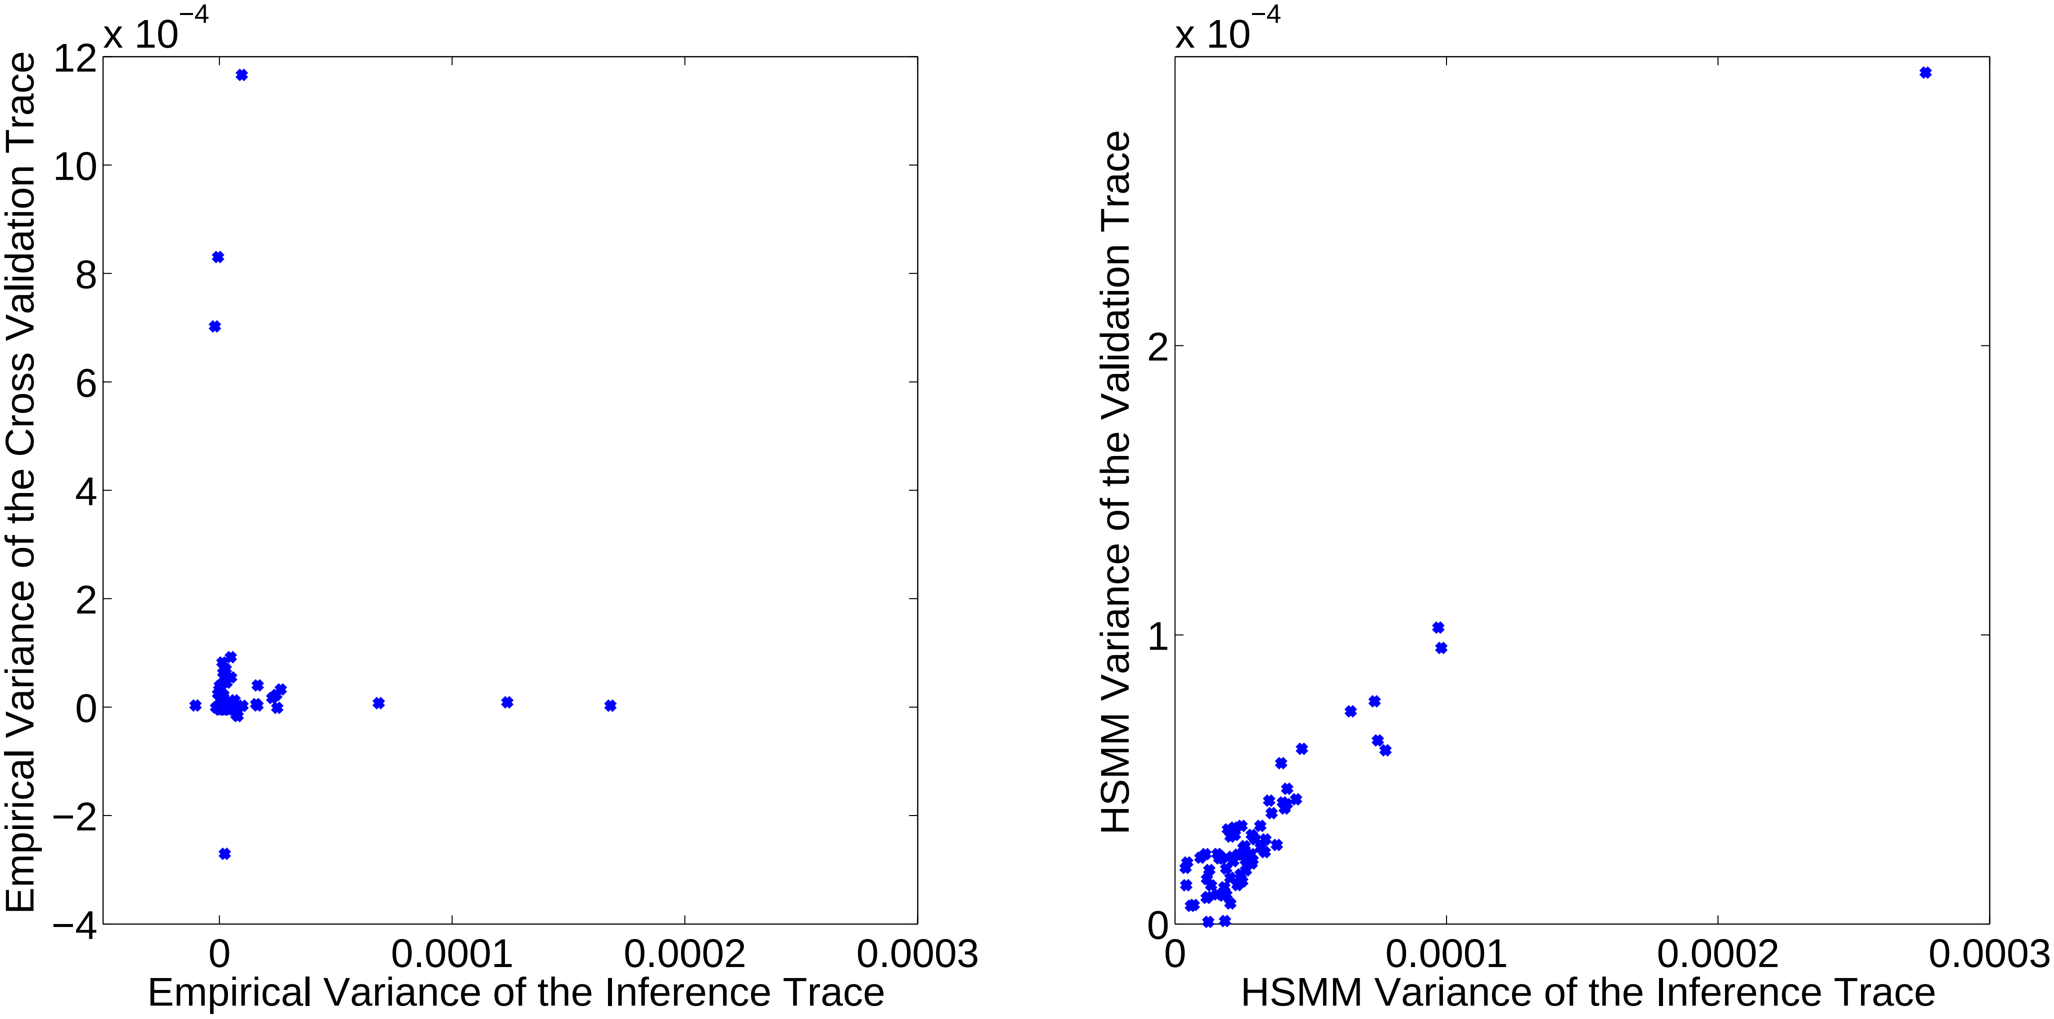

Validation: confidence intervals

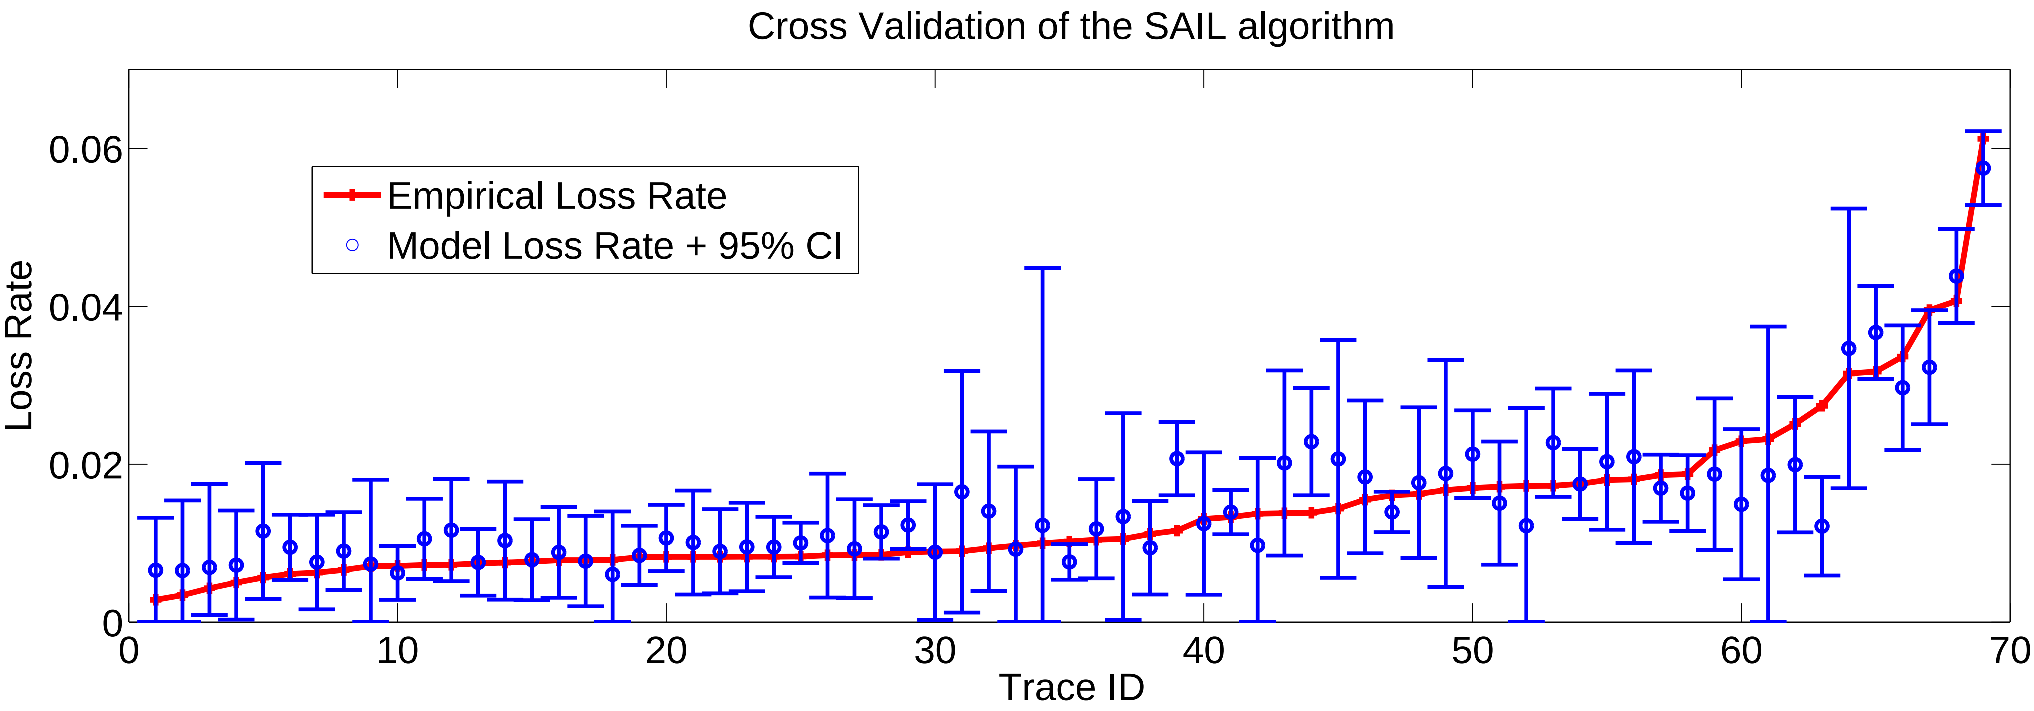

- 95th percentile confidence intervals match in 92% of 1,100 traces
  - independent CIs given 63% coverage
- width of CIs is LARGE

Validation: Web data

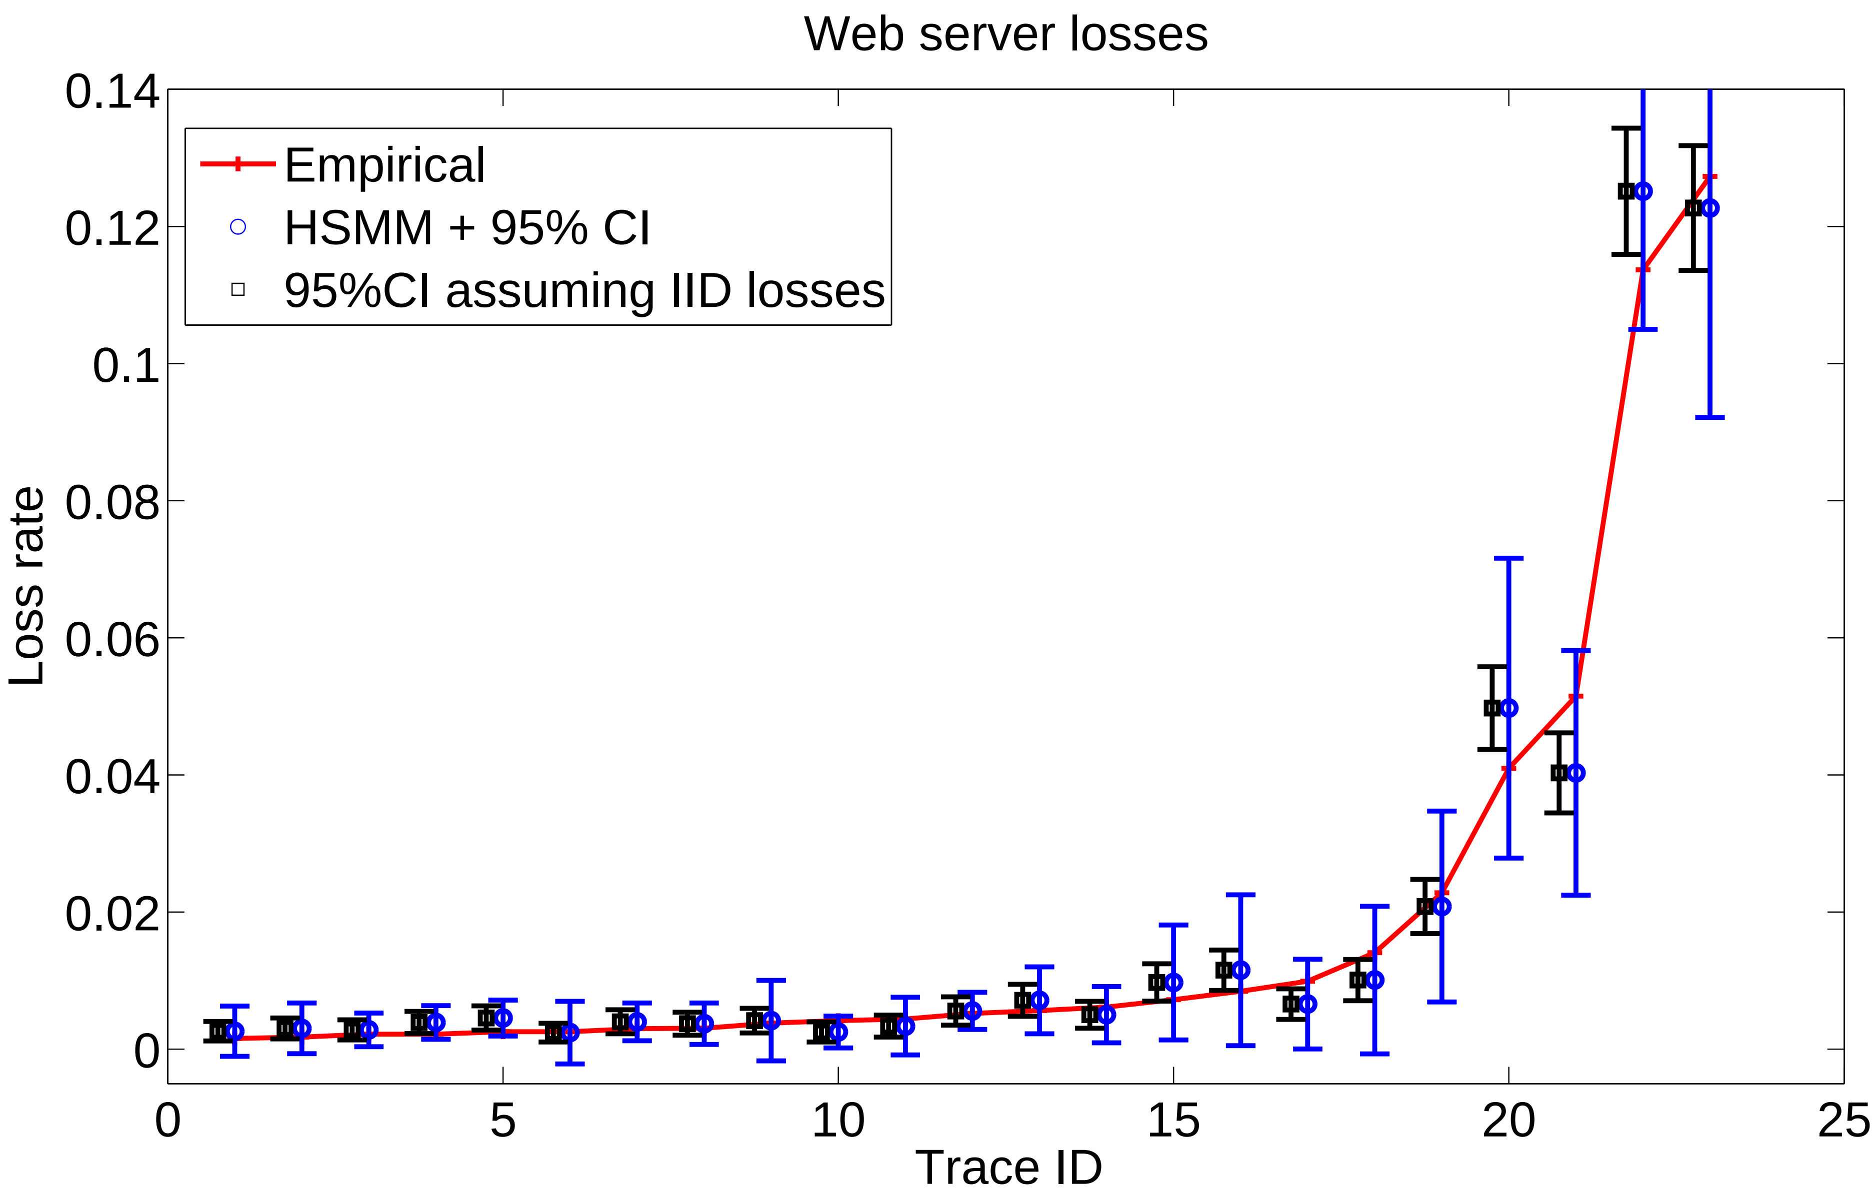

Validation: DNS data

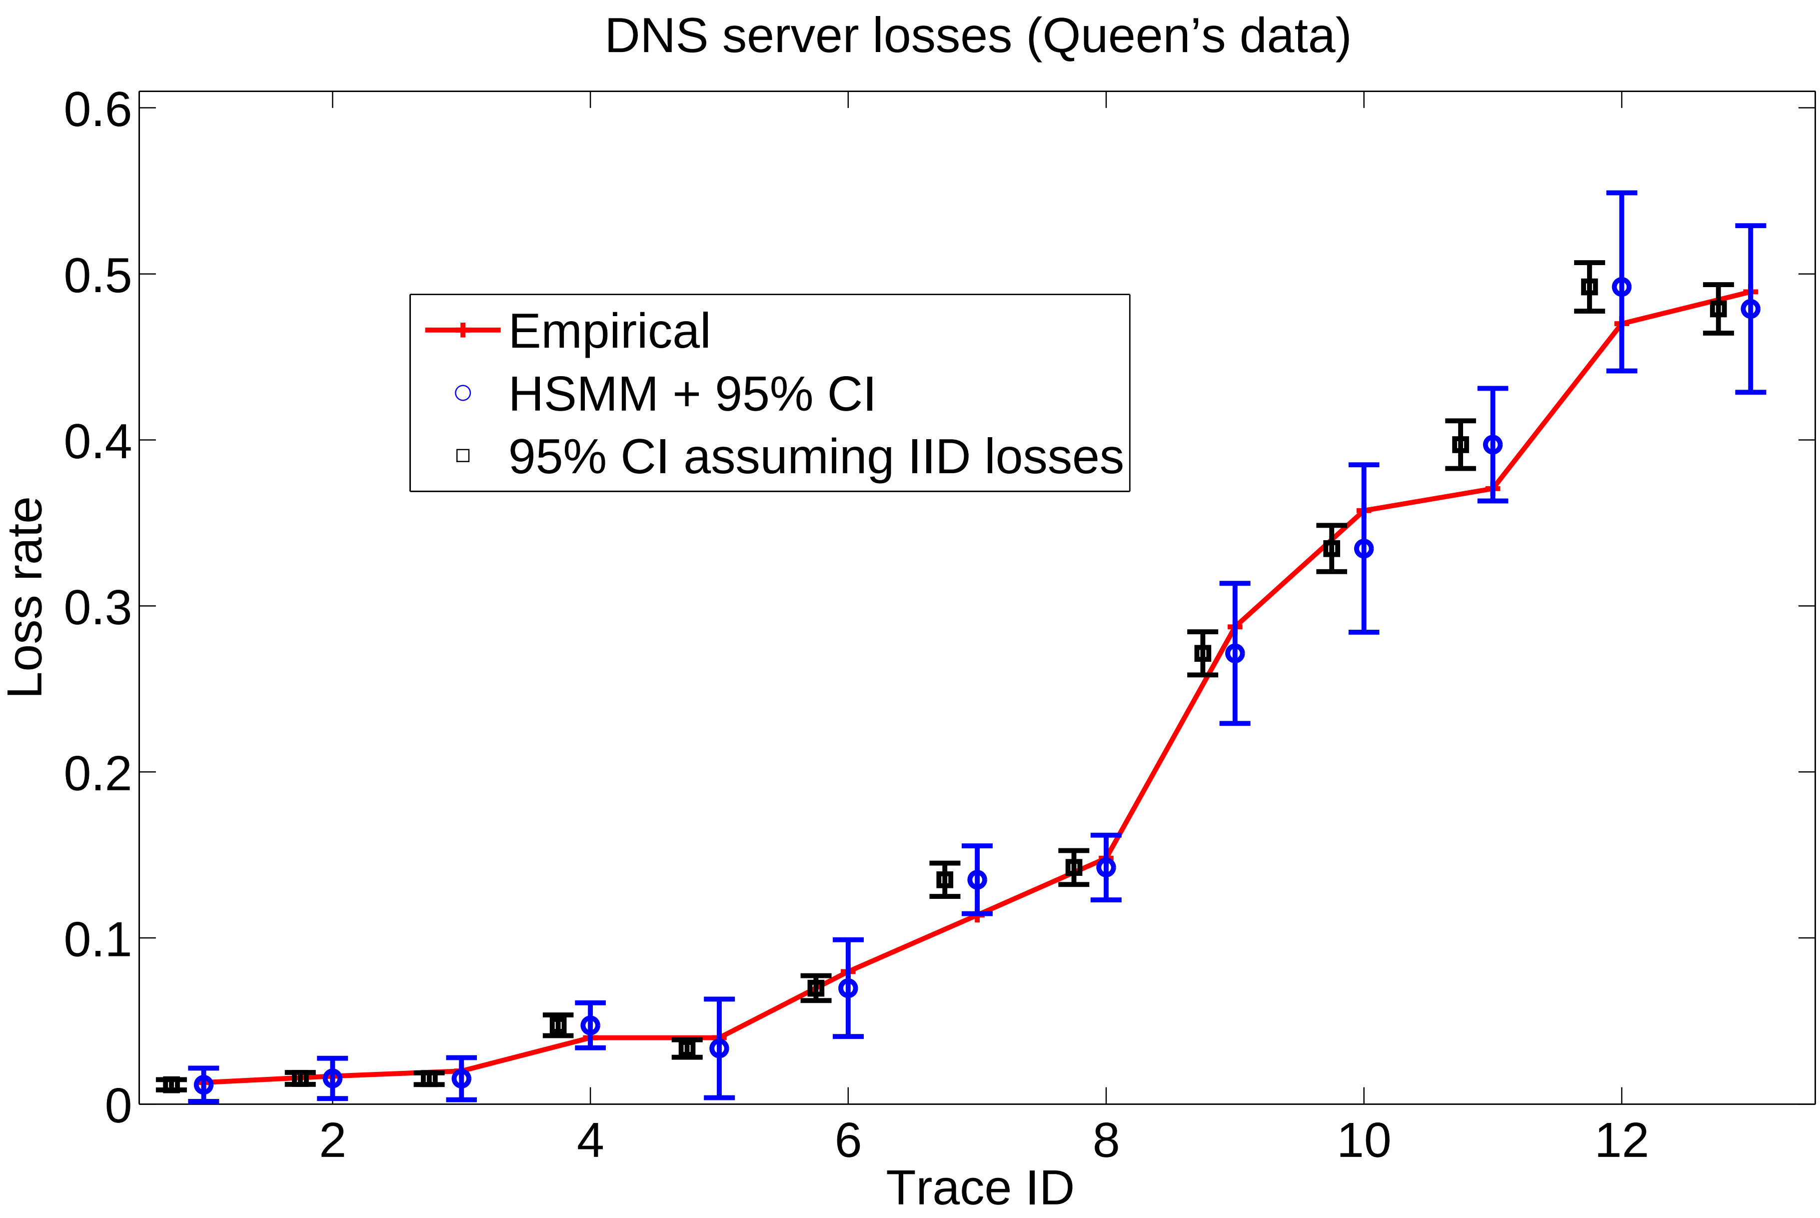

Conclusion

- Measurements needs stats
  - How do you know when a measurement is any good otherwise?
  - Need to deal with all the difficulties of Internet
    - e.g., correlations
- What if I measure every packet?
  - measurements today are for prediction tomorrow
- Loss measurements in particular
  - we have a working technique for loss measurements that copes with the correlations
- Full technical report, data and code are available at `http://www.adelaide.edu.au/directory/hung.nguyen`.

Accuracy

What do we mean by accuracy?

- probes are samples from some underlying process
  - assume measurement equipment has no flaws
- measurements over time can be averaged to give estimates
- implicit assumption of stationary ergodic process
  $\Rightarrow$ a time average converges to an ensemble average
- variance can be estimated by the **Central Limit Theorem**

Accuracy of **estimates** not individual measurements

Central Limit Theorem

For a set of independent, identically distributed RVs X_i, the sample mean

$$\hat{X} = \frac{1}{N}\sum_{i=1}^{N} X_i,$$

has expectation $E\left[\hat{X}\right] = E\left[X_1\right]$, and as $N \to \infty$

$$\sqrt{N}\left(\hat{X} - E\left[X_1\right]\right) \to N(0, \sigma^2)$$

in distribution, where $\sigma^2 = \mathrm{Var}\left[X_1\right]$.

So the 95th% CIs of estimate $\hat{X}$ are $\pm 1.96\sigma/\sqrt{N}$.

Example

White noise

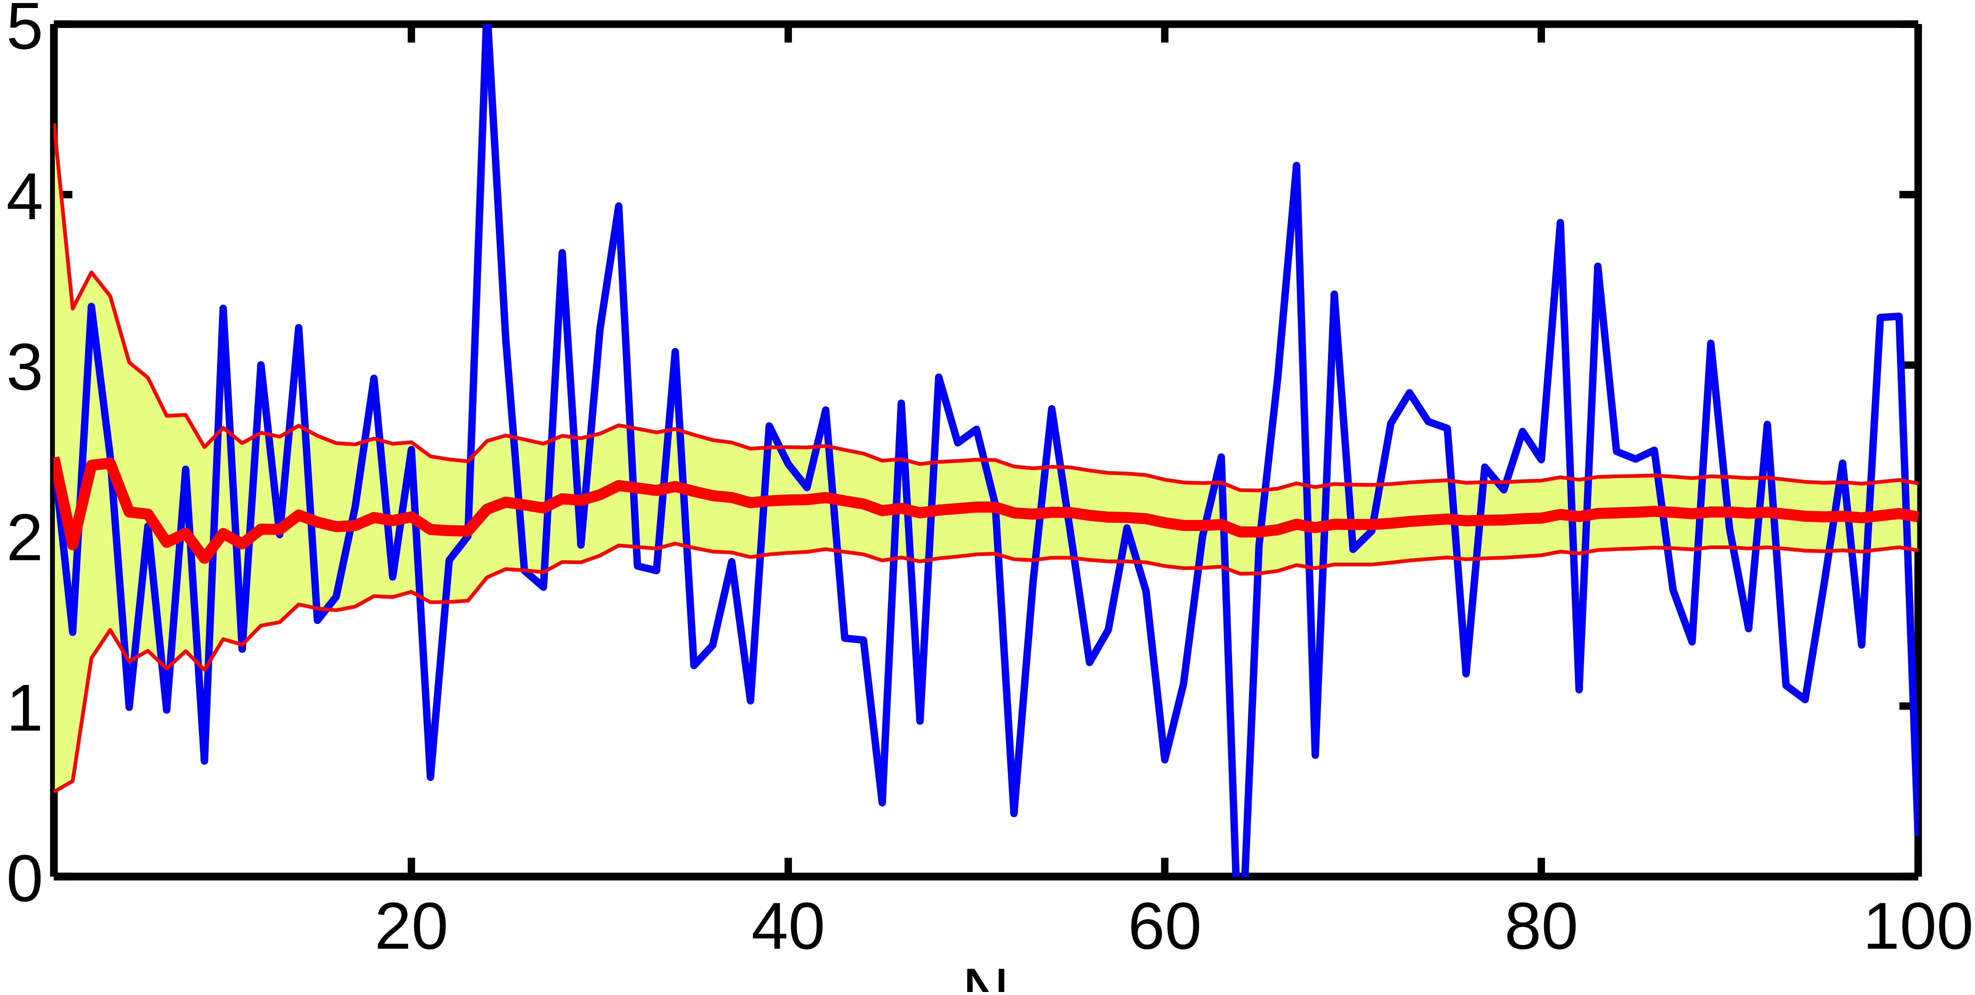

CLT: continuous time

Continuous time process $X(t)$ where the sample mean

$$\hat{X} = \frac{1}{T}\int_0^T X(u)du$$

converges to the true mean $\hat{X} \to E[X]$, and

$$\sqrt{T}\left(\hat{X} - E[X]\right) \to N(0, s^2)$$

in distribution as $T \to \infty$, where

$$s^2 = 2\sigma^2 \int_0^\infty r(u)\, du$$

where $\sigma^2 = \mathrm{Var}[X]$, and $r(s)$ is the autocorrelation of $X(t)$.

Discrete samples

- Correlations are not only a continuous time problem
- Discrete (uniform) samples (interval δt)

$$s^2 = \sigma^2 \left[1 + 2 \sum_{i=1}^{\infty} r(k\,\delta t) \right]$$

- Poisson samples (rate λ)

$$s^2 = \sigma^2 \left[\frac{1}{\lambda} + 2 \int_0^{\infty} r(u)\, du \right]$$

A simple example: the M/M/1 queue

- Poisson packet arrivals (rate λ)
- Exponential service times (mean $1/\mu$)
- Q = number in system (for $\rho = \lambda/\mu$)

$$E[Q] = \frac{\rho}{1-\rho} \quad \text{and} \quad \mathrm{Var}[Q] = \frac{\rho}{(1-\rho)^2}$$

- Autocorrelation (Whitt 1989 [17], Morse 1955 [18])

$$r(t) = \frac{\lambda^2}{(\mu-\lambda)^2} + (\mu-\lambda)\frac{\lambda\mu}{\pi}\int_0^{2\pi} \frac{\sin^2\theta\, e^{-t(\mu+\lambda-2\sqrt{\lambda\mu}\cos\theta)}}{(\mu+\lambda-2\sqrt{\lambda\mu}\cos\theta)^3}\, d\theta$$

- asymptotic variance for M/M/1 (Whitt 1989)

$$s^2 \simeq \frac{4\rho^2}{(1-\rho)^4}$$

Results M/M/1

$\rho = 0.7$

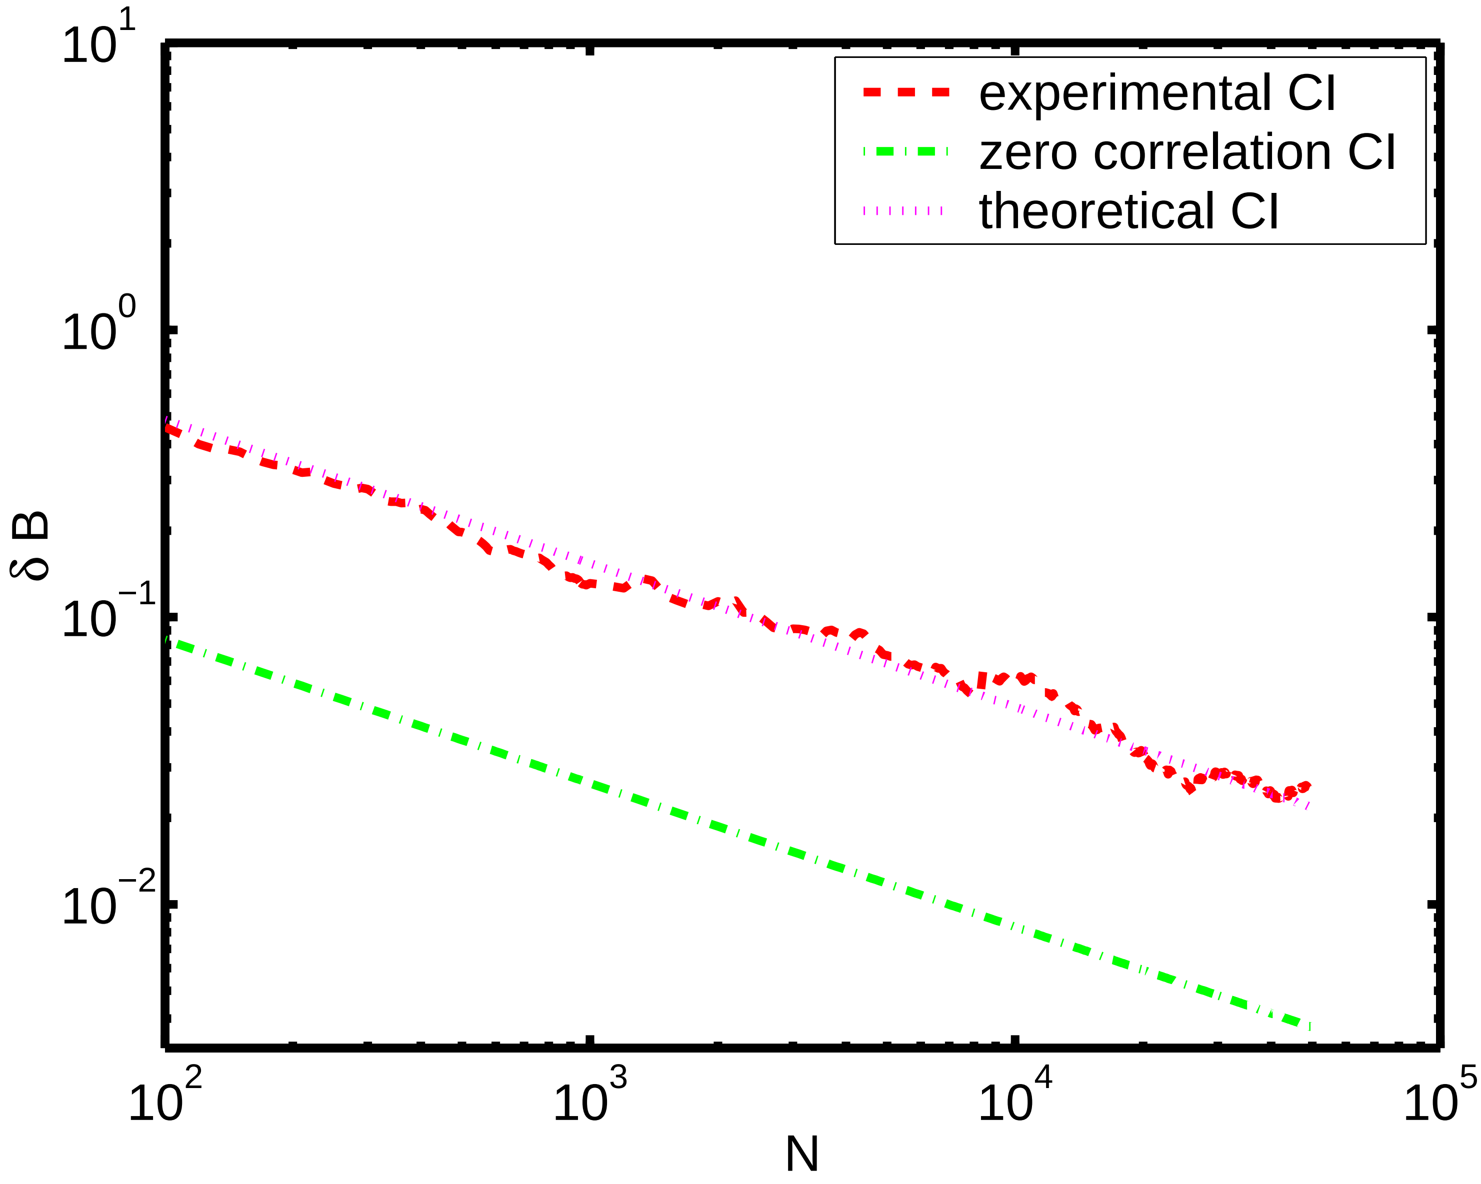

Data

Validity of Semi-Markov Model

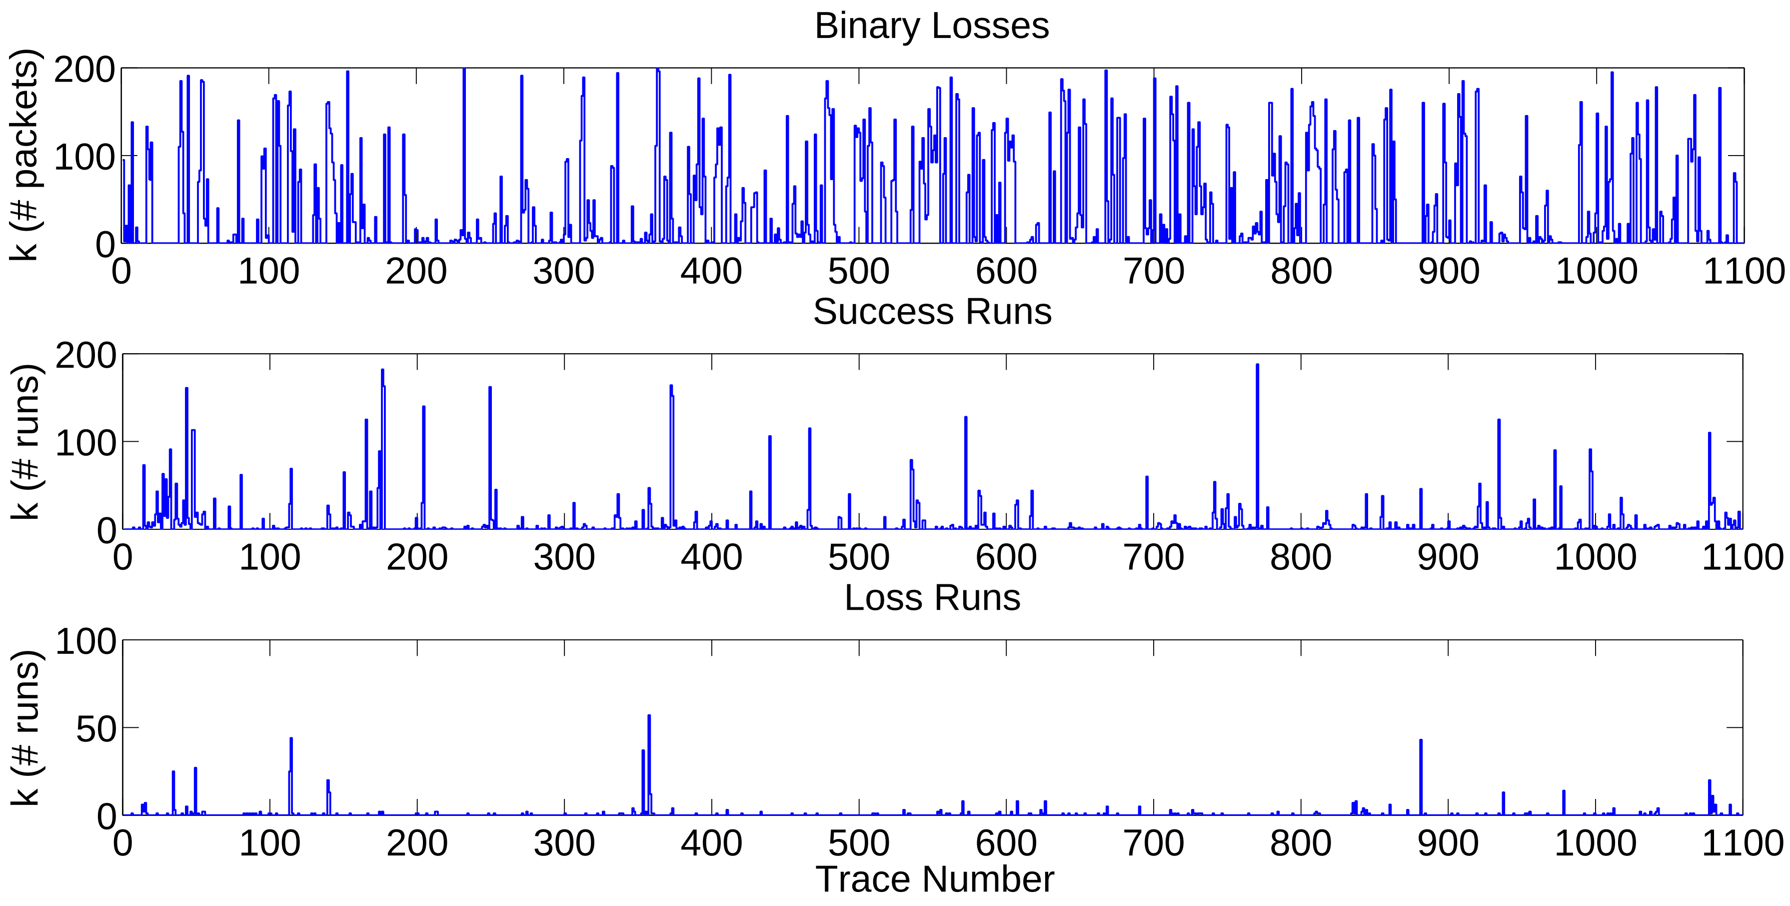

Simulations: comparisons

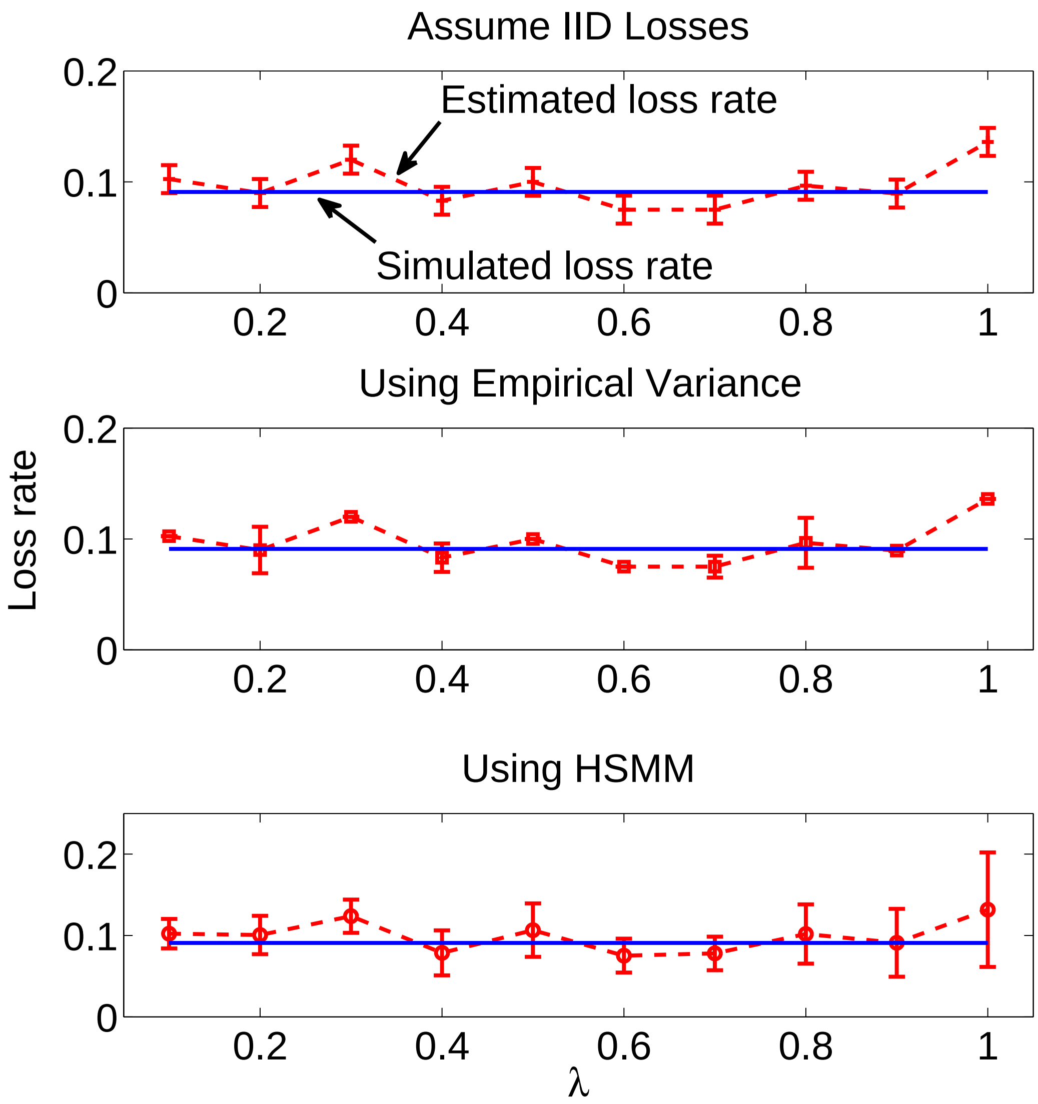

References

[1] Y. Vardi, "Network tomography," Journal of the American Statistical Association, March 1996.

[2] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in ACM SIGMETRICS, (San Diego, California), pp. 206–217, June 2003.

[3] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "An information-theoretic approach to traffic matrix estimation," in ACM SIGCOMM, (Karlsruhe, Germany), pp. 301–312, August 2003.

[4] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "Estimating point-to-point and point-to-multipoint traffic matrices: An information-theoretic approach," IEEE/ACM Transactions on Networking, vol. 13, pp. 947–960, October 2005.

[5] Y. Zhang, M. Roughan, W. Willinger, and L. Qui, "Spatio-temporal compressive sensing and Internet traffic matrices," in ACM Sigcomm, (Barcellona, Spain), pp. 267–278, August 2009.

[6] M. Roughan, "A case-study of the accuracy of SNMP measurements," Journal of Electrical and Computer Engineering, vol. 2010, 2010. Article ID 812979, doi:10.1155/2010/812979. http://www.hindawi.com/journals/jece/2010/812979.html.

[7] L. Li, D. Alderson, W. Willinger, and J. Doyle, "A first-principles approach to