Improving Accuracy in End-to-end Packet Loss Measurement

Joel Sommers, Paul Barford, Nick Duffield,[*] Amos Ron

University of Wisconsin-Madison
AT&T Labs-Research[*]

Background

- Understanding its basic characteristics is important
  - Transport protocol design, throughput modeling, overlay monitoring and optimization
- Standard ways to measure packet loss
  - Passive (SNMP, tcpdump)
  - Active (ping, Poisson modulated probes)

Loss characteristics of interest

loss episode frequency
(fraction of time queue is congested):
((b-a) + (d-c)) / T

mean loss episode
duration:
((b-a) + (d-c)) / 2

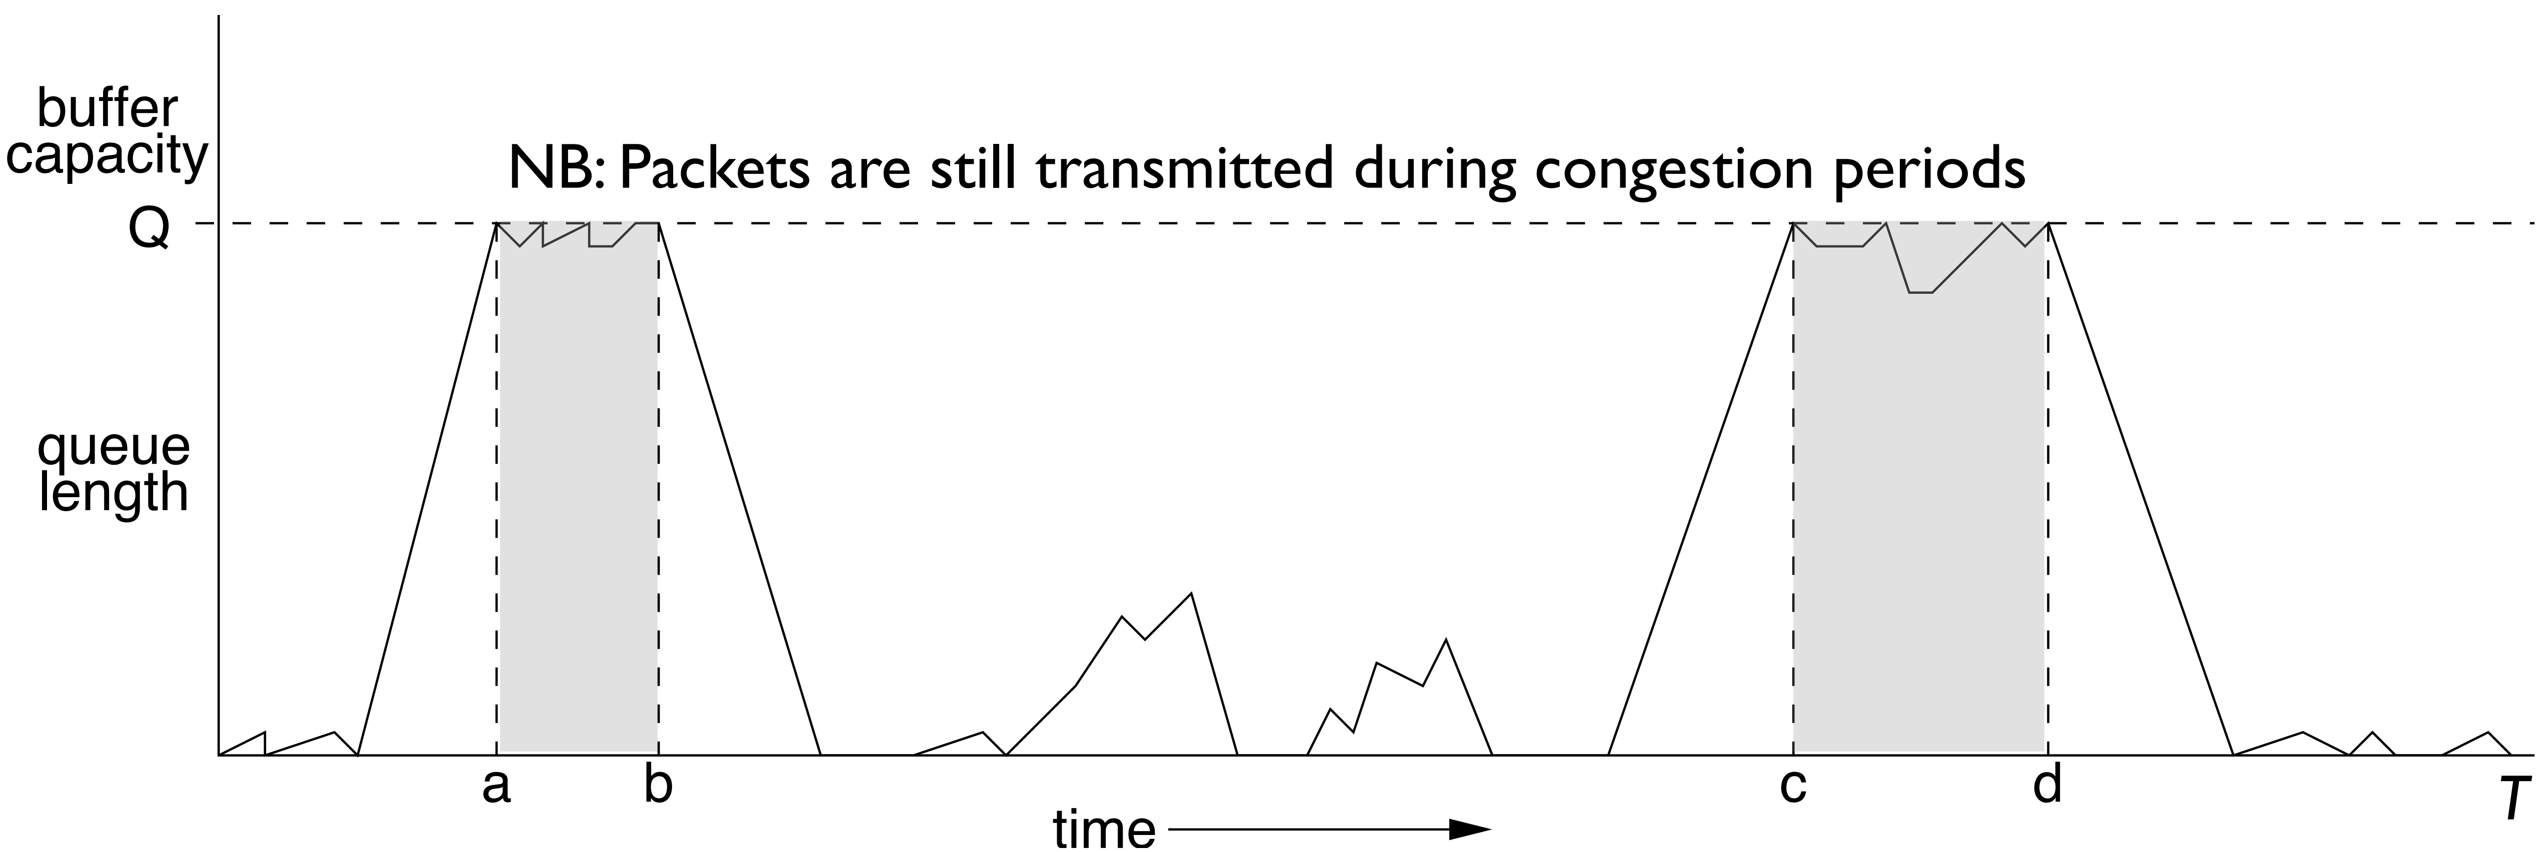

Focus of our study

- How well does traditional Poisson sampling work?
  - What are its limitations? What can be done better?
- Design new sampling process
  - Theory and heuristics
- Controlled laboratory evaluation
  - Compare with Poisson sampling

How well does traditional Poisson sampling work?

- Evaluate frequency and duration estimates
  - Controlled laboratory setting
  - Three kinds of cross traffic
  - Probe rates and packet sizes as [ZPDS01]
  - Experiment duration (15 min) should allow frequency estimates to be close to true frequency

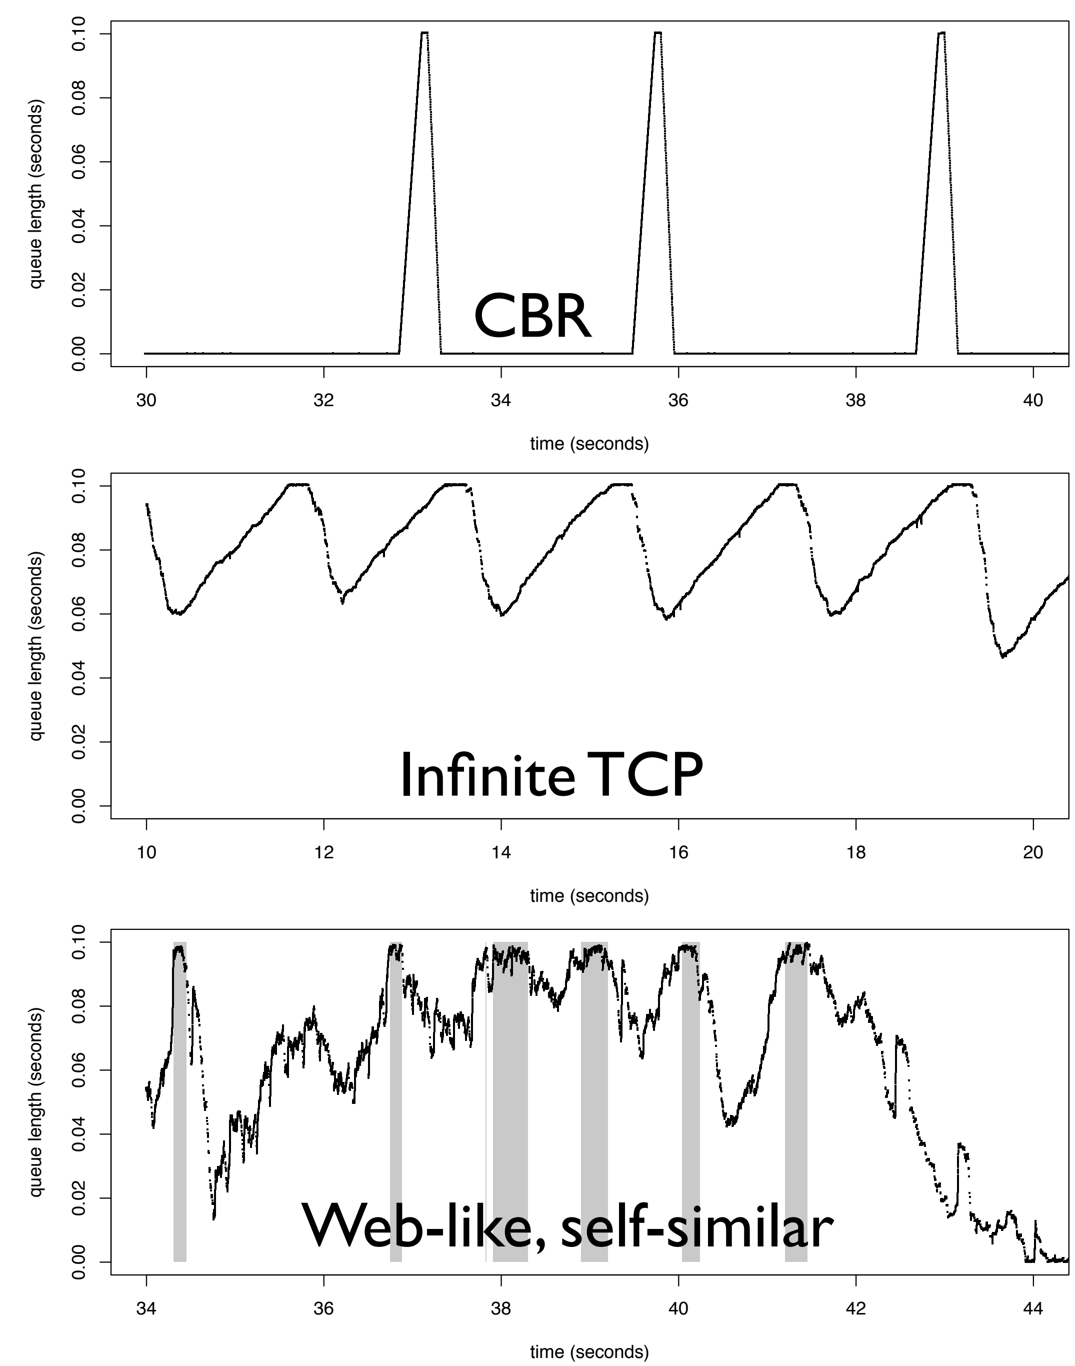

Evaluation of traditional Poisson sampling

- CBR
  - Frequency estimate off by 40%
    Duration estimate off by 85%
- Infinite TCP
  - Very poor frequency estimates
    Duration estimates are 0
- Web-like (table to right)

	frequency	duration (sec)
true values	0.0093	0.136
Poisson (10 Hz)	0.0014	0.000
Poisson (20 Hz)	0.0012	0.022

Lessons and hypotheses

- Poisson sampling is relatively ineffective for estimating congestion frequency and duration
  - ➡ use multi-packet probes
    - Single packet probes often do not experience loss episodes
  - ➡ use loss and delay correlation heuristics
  - ➡ create sampling process to improve duration estimates

Multi-packet probes

- Single packet miss congestion episodes
- Probes with a few packets are more likely to see congestion episodes
- Too many probes distort measurements

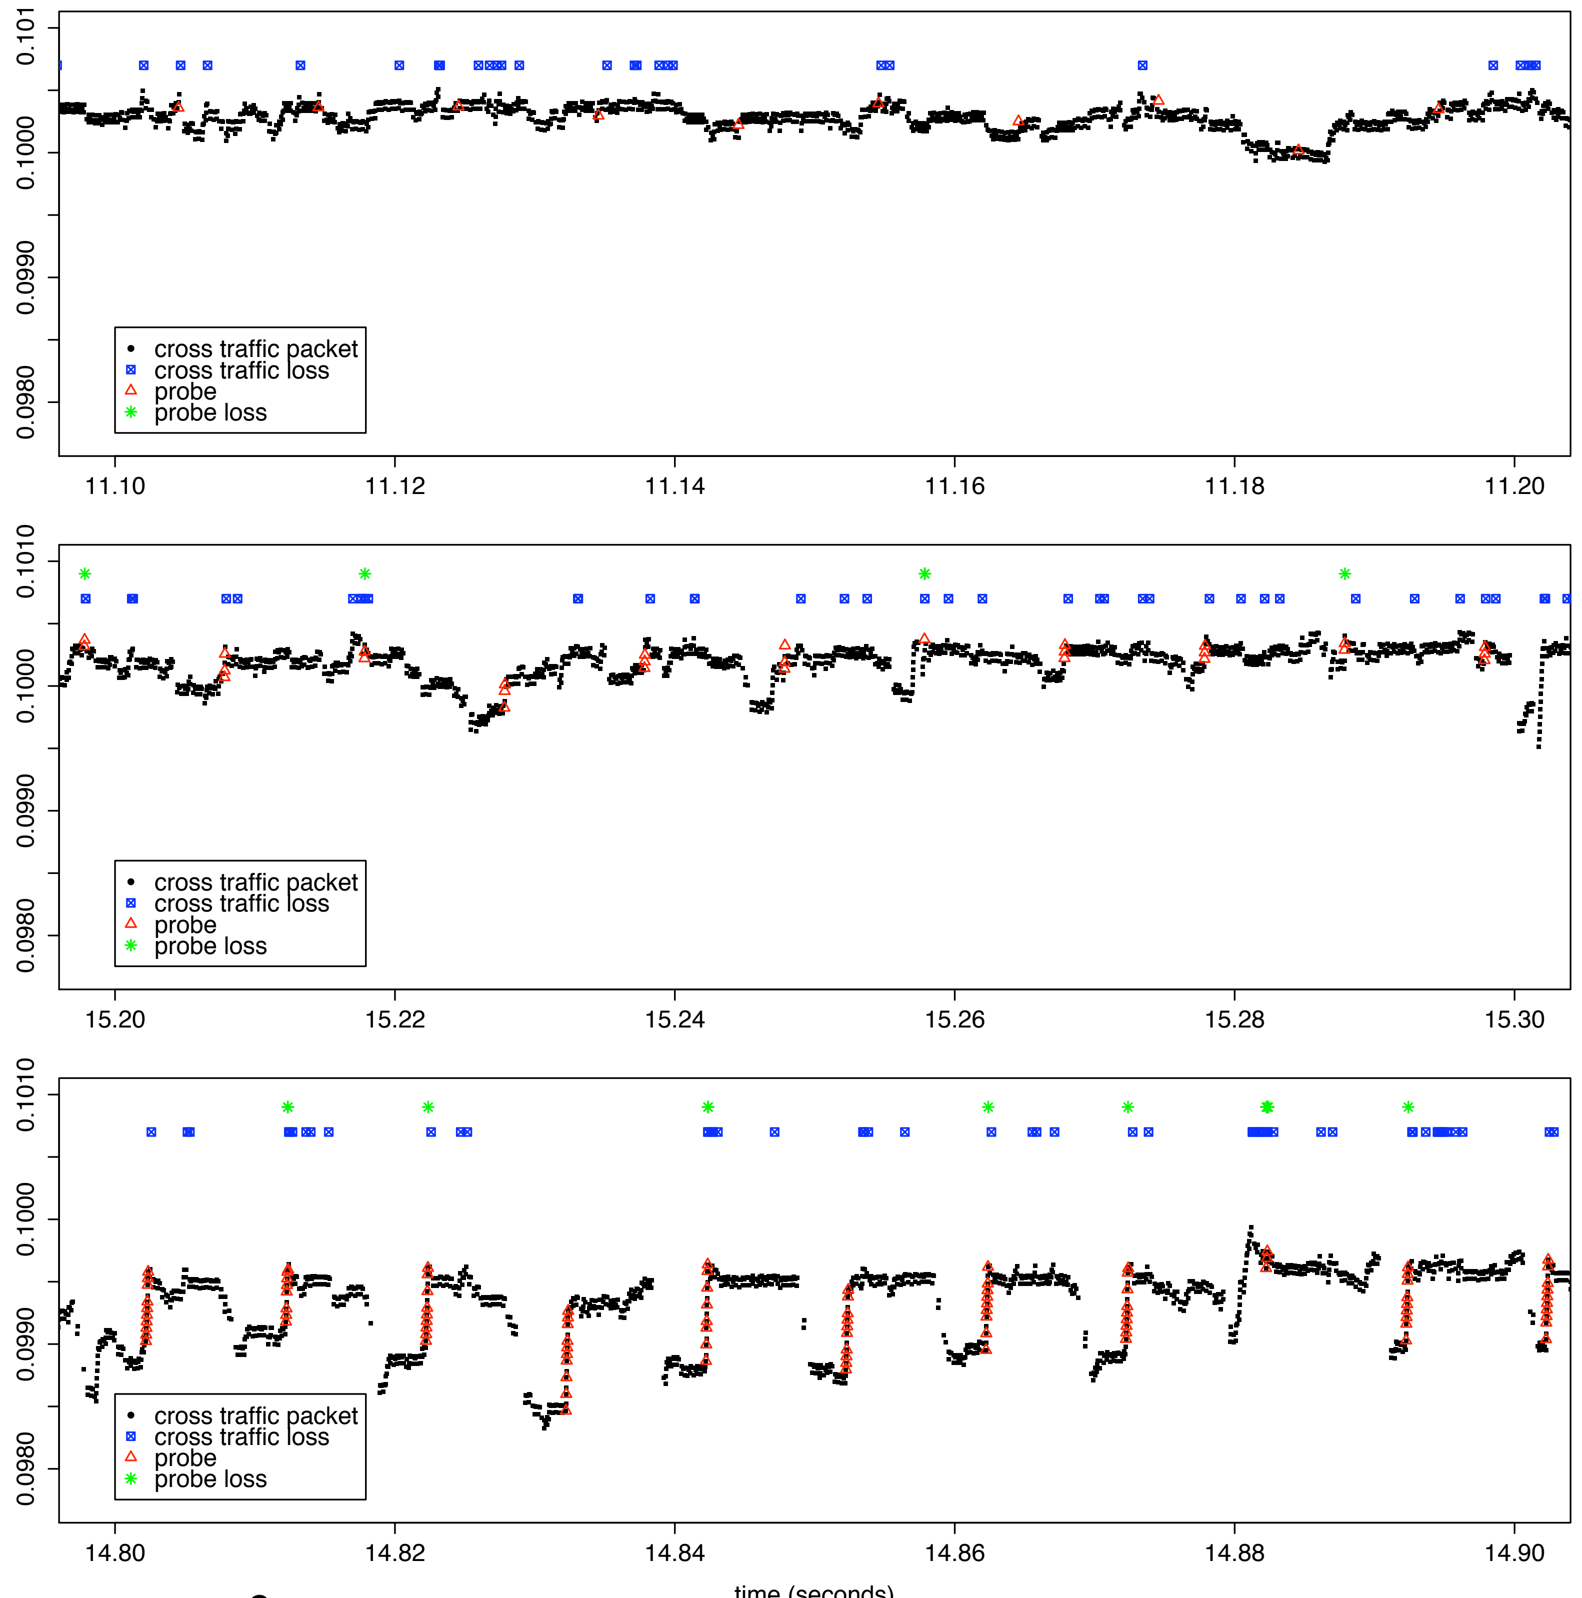

Probe process model

- At the sender
  - Send two multi-packet (3) probes in succession, initiated with probability r at discrete time slot i
  - Individual probe gives instantaneous measure of congestion
  - Probe pairs used to determine congestion dynamics
- At the receiver
  - Record time slots as congested (1) or uncongested (0), using actual packet loss and one-way delay heuristics
  - y_i records congestion as two-digit binary number
  - Y_i denotes true congestion along the path

Key assumptions

- Assume probes don't lie ... usually
  - If there is truly congestion (Y_i), the probes see the effect
  - If y_i is incorrect, assume it is a false negative (y_i = 00)
  - y_i equals Y_i with probability p_k, which is independent of *i* and depends only on the number *k* of 1-digits in Y_i
- For basic algorithm, assume
  - $p_{\{01,10\}} = p_{\{11\}}$ for consistent estimation of duration
  - $p_{\{01,10\}} = p_{\{11\}} = 1$ for consistent and unbiased frequency estimation

One-way delay and congestion heuristics

- Improve single probe measurement of congestion
  - Probes within τ seconds of true loss $\Rightarrow$ congestion
  - Probes with OWD $\geq (1-\alpha)$ OWDmax $\Rightarrow$ congestion
- Observations from sensitivity experiments
  - Relationship between larger parameter value and more congestion inferred
  - Tradeoff between probe rate and parameter settings

New probe model example

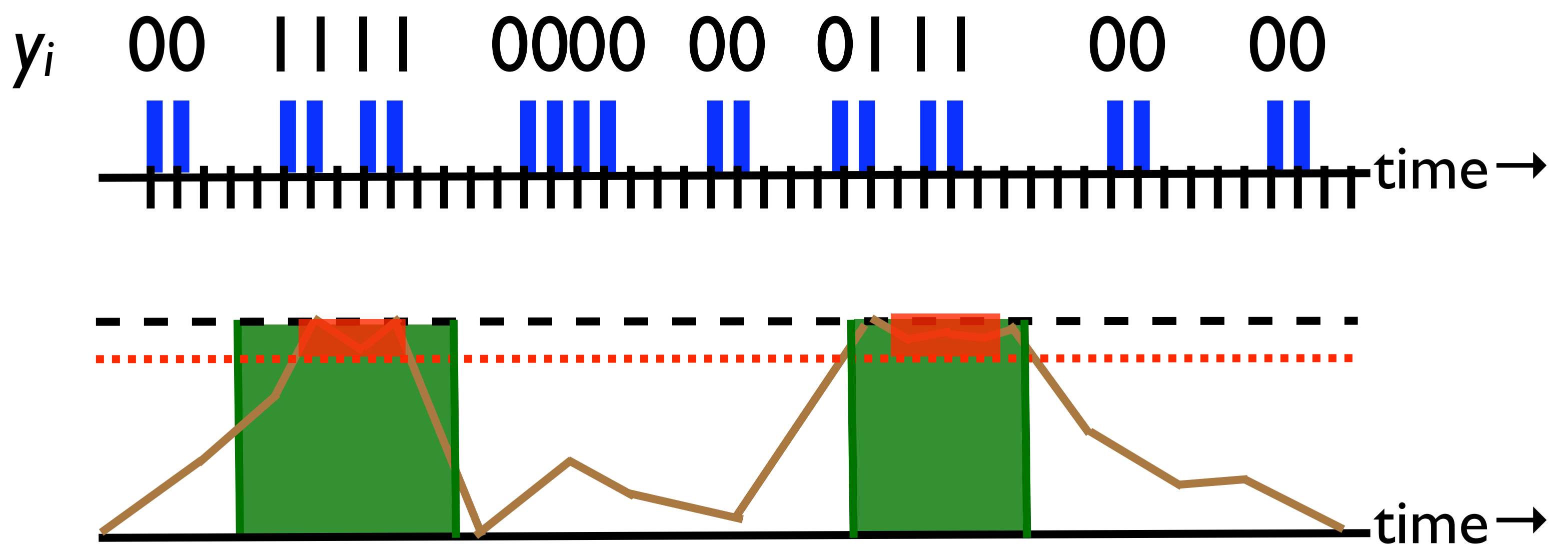

Red line denotes α OWD threshold heuristic

Green areas denote τ loss proximity heuristic

Estimating congestion frequency

$$\hat{F} = \sum_i z_i / M$$

- z_i is a random variable whose value is the first digit of y_i
- M is the total number of probe pairs
- Estimator is unbiased, and under mild conditions, consistent

Estimating congestion duration (1)

- Assume we have knowledge of the path at all possible time slots in our discretization
  - For k=1,2,..., there were exactly j_k congestion episodes of length k
    - Congestion occurred over total of A time slots, $A = \sum k j_k$
    - Total number of congestion episodes is $B = \sum j_k$
    - Average duration D of a congestion episode is therefore $D := A/B$

Estimating congestion duration (2)

Note that there are B time slots i for which $Y_i = 01$,
and also B time slots i for which $Y_i = 10$

Note also that there are exactly $A+B$ time slots i for which $Y_i \neq 00$

Define $R := \#\{i : y_i \in \{01,10,11\}\}$ and $S := \#\{i : y_i \in \{01,10\}\}$

We arrive at $\mathsf{E}(R)/\mathsf{E}(S) = \frac{p_2(A-B) + 2p_1 B}{2p_1 B}$

Assuming $p_{\{01,10\}} = p_{\{11\}}$, the estimator for the mean congestion duration is therefore

$$\hat{D} := 2 \times \frac{R}{S} - 1$$

Validation of output

- Monitor results in real-time to check whether assumptions have been violated and to increase confidence in results
  - Probability of $y_i = 01$ is assumed to be same as $y_i = 10$ — monitor these rates of occurrence
  - $p_{\{01,10\}} = p_{\{11\}}$ for consistent estimation of duration
  - $p_{\{01,10\}} = p_{\{11\}} = 1$ for consistent and unbiased frequency estimation

Laboratory results summary

- Implemented new sampling model in a tool called badabing
- Experiments in a controlled testbed using a range of probe rates and range of thresholds for inferring congestion
  - Estimates are often within 25% of actual congestion frequency and duration values; many within 10%
  - A significant improvement over traditional Poisson sampling for both frequency and duration estimation

badabing evaluation (CBR, single episode type)

	loss frequency		loss duration	
r	true	badabing	true	badabing
0.1	0.0069	0.0016	0.068	0.054
0.3	**0.0069**	**0.0065**	**0.068**	**0.073**
0.5	0.0069	0.0060	0.068	0.051
0.7	0.0069	0.0070	0.068	0.051
0.9	0.0069	0.0078	0.068	0.053

badabing evaluation (web-like, self-similar traffic)

	loss frequency		loss duration	
r	true	badabing	true	badabing
0.1	0.0044	0.0017	0.060	0.071
0.3	**0.0011**	**0.0011**	**0.113**	**0.143**
0.5	0.0114	0.0117	0.079	0.074
0.7	0.0043	0.0039	0.071	0.076
0.9	0.0031	0.0038	0.073	0.062

Comparing badabing with Poisson probes

- With same probe stream rate for Poisson and badabing
  - Constant bit rate cross traffic
    - Both frequency and duration estimates are within 7% for badabing; Frequency estimate off by 40% and duration estimate off by 85% for Poisson
  - Web-like cross traffic
    - Badabing correctly estimates frequency and duration estimate is within 25%; Each estimate derived from Poisson-modulated probes is at least 80% off

Summary

- Simple Poisson sampling is relatively ineffective for measuring congestion frequency and duration
- Badabing provides more accurate estimation of congestion frequency and duration
    - Estimator performance depends only on total number of probes sent, not on sending rate
    - Simple validation methods for measurement output
    - Accuracy improvements (and basic assumptions) validated in a laboratory testbed

the end

http://wail.cs.wisc.edu/