Improving Accuracy in End-to-end Packet Loss Measurement

Paper by Joel Sommers, Paul Barford, Nick Duffield, and Amos Ron

Presented by Christoph Jechlitschek

Motivation

- Understanding the impact of packet loss
  - New Reno and Sack
- TCP throughput modeling
- Infer network properties from losses
  - Cross traffic intensity
  - Bottleneck buffer size

2

Packet loss measurements

- Passive
  - SNMP
  - Tcpdump
- Active
  - Ping
  - Poisson modulated probes

3

Simple Queue model

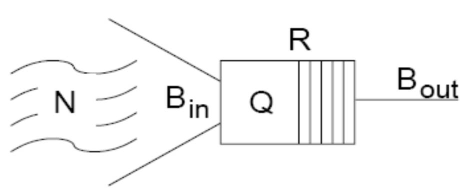

figure from [1]

4

Loss characteristics

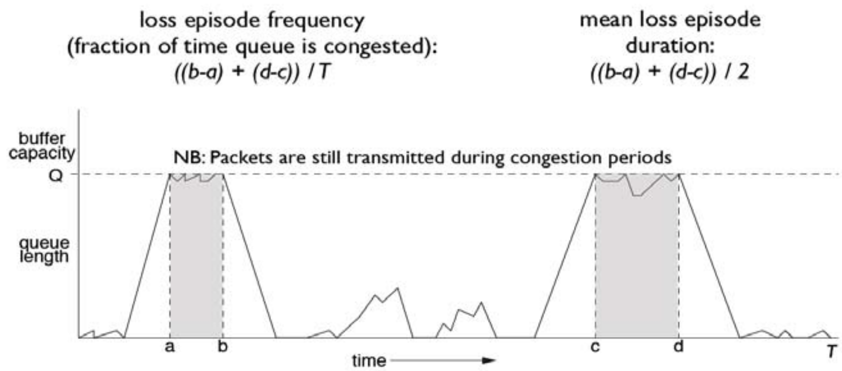

figure from [2] 5

Test environment

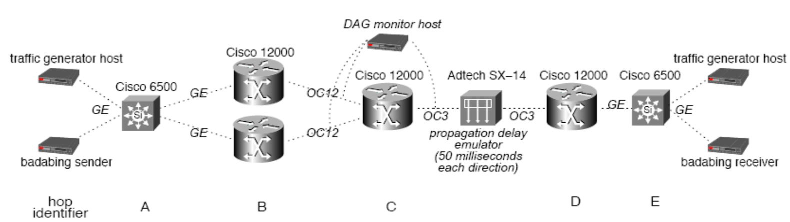

figure from [1] 6

The Competitor: Zing

- Sends UDP probes at Poisson-modulated intervals
- Uses the “well known” PASTA principle
- Drawbacks:
  - Misses losses
  - Needs longer measurement periods

7

Performance of Zing

	frequency	duration μ (σ) (seconds)
true values	0.0265	0.136 (0.009)
ZING (10Hz)	0.0005	0 (0)
ZING (20Hz)	0.0002	0 (0)

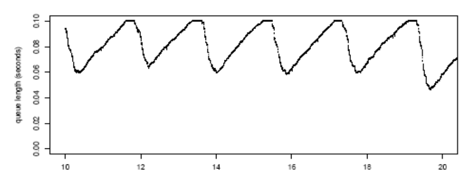

	frequency	duration μ (σ) (seconds)
true values	0.0069	0.068 (0.000)
ZING (10Hz)	0.0036	0.043 (0.001)
ZING (20Hz)	0.0031	0.050 (0.002)

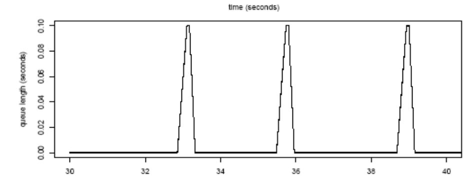

	frequency	duration μ (σ) (seconds)
true values	0.0093	0.136 (0.009)
ZING (10Hz)	0.0014	0 (0)
ZING (20Hz)	0.0012	0.022 (0.001)

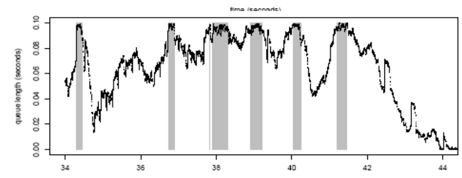

figures and tables from [1]

Badabing

- Divide time into discrete time slots
- With some fixed prob. p commence experiment in each slot
- Experiment = 2 probes at current and next slot
- Record outcome as two digit number e.g. y_i = 00 or y_i = 01
- A congestion has the form 01…10

9

Badabing

- Congestion frequency: $\hat{F} = \sum_k z_i / M$
- Congestion duration: $\hat{D} = 2 \times \frac{R}{S} - 1$

10

Congestion Duration estimation

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

slide content from [2] 11

Congestion Duration estimation

- 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 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 $E(R)/E(S) = (p_2(A-B) + 2p_1B)/2p_1B$
- Assuming $p_{\{01, 10\}} = p_{\{11\}}$, the estimator for the mean congestion duration is therefore:

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

slide content from [2] 12

Example Basic Algorithm

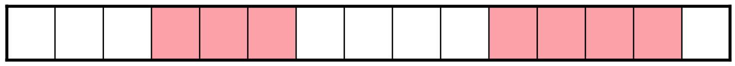

Y_i	00	00	01	11	11	10	00	00	00	01	11	11	11	10	00
y_i	00			11		10	00		00		11	11		10	

Time slots N = 15, $p_1 = p_2 = 1$

Y_i	$\hat{F} = 7/15 = 0.467$	$\hat{D} = 2\frac{9}{4} - 1 = 3.5$
y_i	$\hat{F} = 5/15 = 0.3$	$\hat{D} = 2\frac{5}{2} - 1 = 4$

13

Improved Algorithm

- Does not assume $p_1 = p_2$
- Need to estimate $r = p_1/p_2$
- Conduct extended experiments – 3 probes
- U:= #{ i: $y_i \in$ {011, 110}}
- V:= #{ i: $y_i \in$ {001, 100}}
- r = V/U

$$\hat{D} := \frac{2V}{U} \times \left(\frac{R}{S} - 1 \right) + 1$$

14

Reliability of Probes

- Not all connections experience loss during congestion
- Send packet trains to increase that chance
- Use delay measurement
  - Probes OWD ≥ $(1-\alpha)OWD_{max} \Rightarrow$ congestion
  - Probes within t seconds of true loss $\Rightarrow$ congestion

15

Reliability of Probes

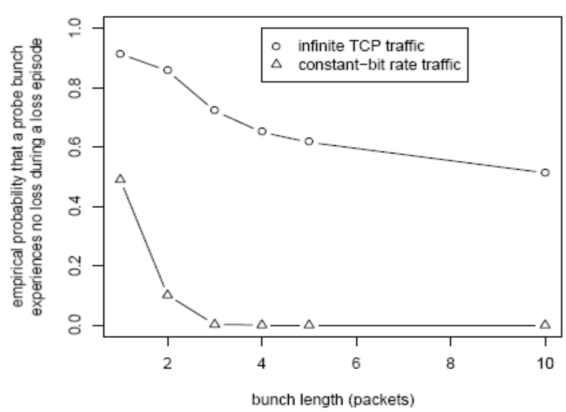

figure from [1]

16

Single probe

Probe train of 3 packets

Probe train of 10 packets

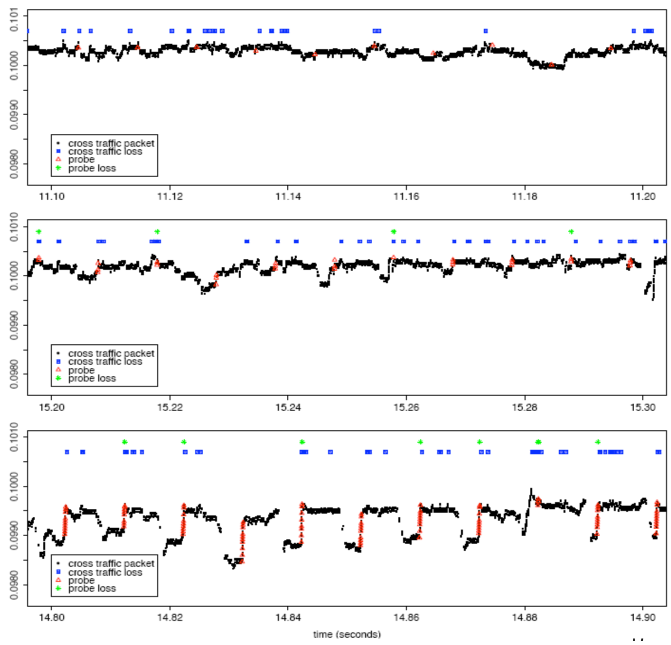

figures from [2]

Delay Heuristic

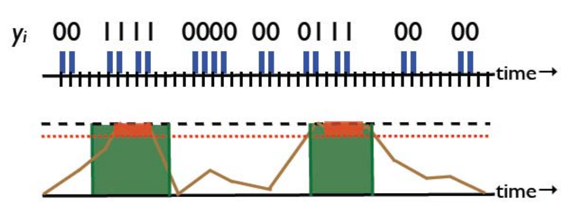

Red line denotes α OWD threshold heuristic

Green areas denote τ loss proximity heuristic

figure from [2]

Results

- Loss estimates for CBR traffic

p	loss frequency		loss duration (seconds)	
	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

- Loss estimates for web-like traffic

p	loss frequency		loss duration (seconds)	
	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

- … but where are the estimates for TCP?

?

tables from [1] 19

Zing vs. Badabing

- With same probe stream rate for Zing and Badabing (876 kb/s)

traffic scenario	tool	loss frequency		loss duration	
		true	measured	true (sec)	measured (sec)
CBR	BADABING	0.0069	0.0065	0.068	0.073
	ZING	0.0069	0.0041	0.068	0.010
Harpoon web-like	BADABING	0.0011	0.0011	0.113	0.143
	ZING	0.0159	0.0019	0.119	0.007

table from [1] 20

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

slide content from [2] 21

Questions

22

References

- [1] Improving Accuracy in End-to-end Packet Loss Measurement, Sigcomm 2005, http://www.acm.org/sigs/sigcomm/sigcomm2005/paper-SomBar.pdf
- [2] Conference slides for Improving Accuracy in End-to-end Packet Loss Measurement, http://www.acm.org/sigs/sigcomm/sigcomm2005/slides-SomBar.pdf

23