Accurate and Efficient SLA Compliance Monitoring

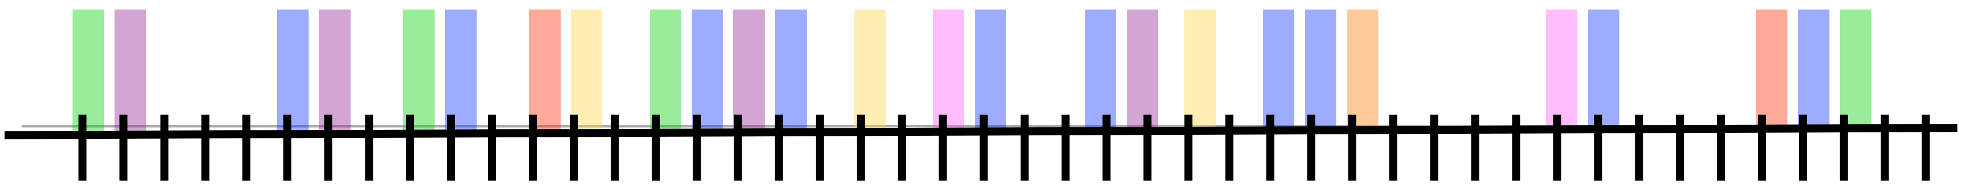

Joel Sommers

Colgate University

Paul Barford

University of Wisconsin-Madison

Nick Duffield

AT&T Labs-Research

Amos Ron

University of Wisconsin-Madison

IETF 71
IPPM Working Group
Philadelphia, PA
10 March 2008

Note: there may be intellectual property associated with some of this material.

Motivation

- Service level agreements (SLAs) specify performance guarantees made by Internet service providers
  - Example metrics: packet loss, delay, delay variation
- Accurate and robust SLA compliance monitoring is important for service providers and their customers
  - Lightweight, effective monitoring is a key challenge
    - Measurement on a single path
    - Network-wide monitoring
  - Non-compliance can have serious consequences!

Overview

- New sampling method for packet loss
- New methods for calculating existing statistics
  - Mean delay, delay percentiles
  - Packet loss average
- New delay variation statistic
- New optimized discrete-time sampling approach
- Evaluation in a controlled laboratory setting
  - Tool (SLAm) accuracy compared with appropriate RFCs

- Sommers, Barford, Duffield, and Ron. "Accurate and Efficient SLA Compliance Monitoring." Proceedings of ACM SIGCOMM, August 2007.

Packet loss

- Geometric sampling
- Builds on Badabing probe methodology [SBDR 05]
  - Each sample consists of two probes, sent in consecutive time slots
    - Each probe defined as three packets sent back-to-back
- New methodology for loss average statistic
  - Use loss episode frequency and mean duration statistics from Badabing
  - Packet loss average is derived from Badabing statistics
- Code available at http://wail.cs.wisc.edu/

One-way delay

- Samples are geometrically distributed
- New methods for calculating statistics
  - Mean delay estimate
    - Based on Simpson's method for numerical integration
  - Delay percentile estimation
    - Statistically sound; does not assume any underlying distribution of delay
    - Result of method is a confidence bound on the desired percentile
  - Inference of delay distribution for an unmeasured (but linearly dependent) path

Delay variation

- Samples are the same as in RFC 3393
  - Periodic samples used in our experiments, but they're not required
- New statistic for estimating DV
  - Closest in spirit to Type-P-One-way-ipdv-jitter statistic in RFC 3393
    - Similar to RTP jitter metric (RFC 3550)
  - Qualitative measure of delay variation along a path
    - Calculated over an ordered set of samples
    - A measure of *distortion* from zero variation

Multi-objective probing

- Sampling approach based on discrete-time clock
  - E.g., for geometric and periodic sampling
  - Probes may be scheduled to be sent at same time slot
  - Tag probes according to the sampling methode to which they apply

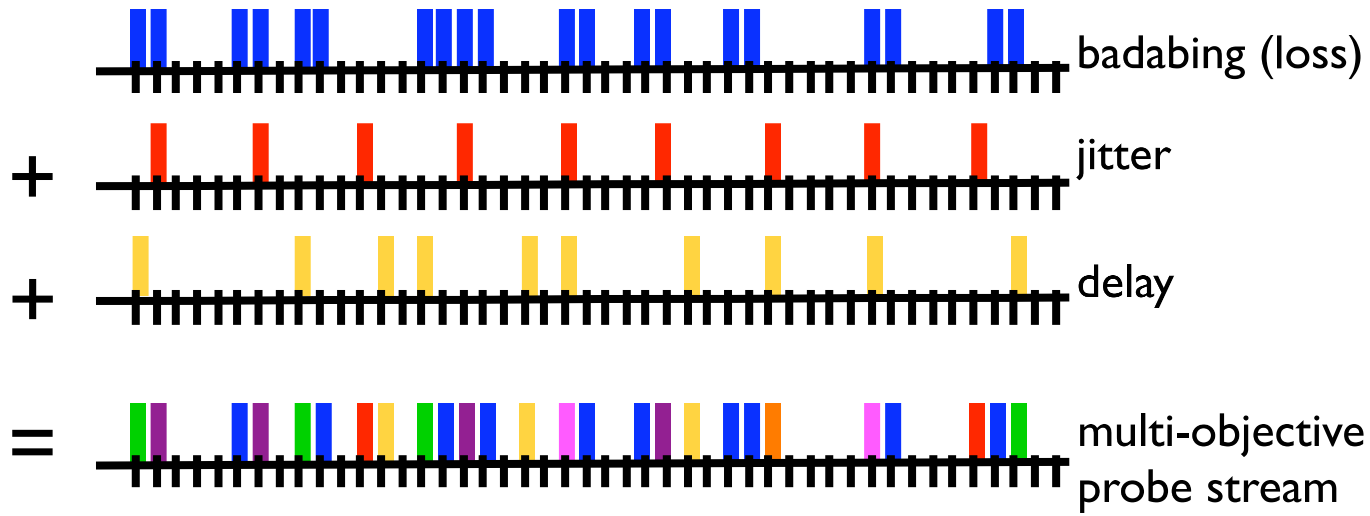

Results

- Evaluated in controlled laboratory environment
  - Two topologies: dumbbell and star
  - A range of background traffic settings and loads
- Compare with RFC-standard probe streams at same bitrate
  - Results for new methods are closer to true values
    - Mean delay results show modest improvement in accuracy
    - Loss average results significantly closer to true values
    - Delay variation statistic is more robust than comparable statistic; more accurately tracks turbulent conditions

Recent Related Work

- B.Y. Choi, S. Moon, R. Cruz, Z.-L. Zhang, C. Diot. Practical delay monitoring for ISPs. ACM CoNext, 2005.
  - Estimating delay percentiles
- F. Baccelli, S. Machiraju, D. Veitch, J. Bolot. The Role of PASTA in Network Measurement. ACM SIGCOMM, 2006.
- F. Baccelli, S. Machiraju, D. Veitch, J. Bolot. On Optimal Probing for Delay and Loss Measurement. ACM IMC, 2007.
  - Identifying unbiased sampling methods with minimal variance
- Y. Chen, D. Bindel, H. Song, R. Katz. An Algebraic Approach to Practical and Scalable Overlay Network Monitoring. ACM SIGCOMM, 2004.
- D.B. Chua, E.D. Kolaczyk, M. Crovella. Efficient Estimation of End-to-end Network Properties. IEEE INFOCOM, 2005.
  - Tomographic inference of performance metrics

Summary

- A set of new methodologies for accurate, lightweight SLA compliance monitoring
  - Multi-objective probing: reduces overhead
  - Delay: accurate estimates of mean and percentiles
  - Loss rate: accurate estimate based on Badabing
  - Delay variation: robust qualitative statistic
- Methodologies implemented in a tool called SLAm
  - Laboratory tests with one- and two-hop topologies
- Paper available at http://cs.colgate.edu/faculty/jsommers/pubs/fp122-sommers.pdf

The End

Questions?

Is the IPPM WG interested in revising/updating existing active probe recommendations?

Evaluate comparative merits of recently proposed measurement techniques?

Results: Delay

- Results for SLAm are closer to true value than standard Poisson-based stream (RFC 2679)
  - Fast convergence to true mean delay (in paper)

Results for self-similar background traffic generated using Harpoon.

mean delay comparison	SLAm		RFC 2679	
	true	estimate	true	estimate
dumbbell (60%)	0.006	0.006	0.007	0.009
dumbbell (75%)	0.014	0.014	0.006	0.013
star: route 1	0.007	0.006	0.007	0.005
star: route 2	0.009	0.008	0.009	0.006
star: route 3	0.005	0.005	0.005	0.004
star: route 4	0.007	0.006	0.007	0.004

Results: Delay Quantiles

- Calculated quantiles with 90% confidence interval
- Intervals generally include true quantile, with few exceptions
  - For all traffic scenarios used, in both dumbbell and star topologies

Results for CBR in star topology (left) and long-lived TCP in dumbbell topology (right)

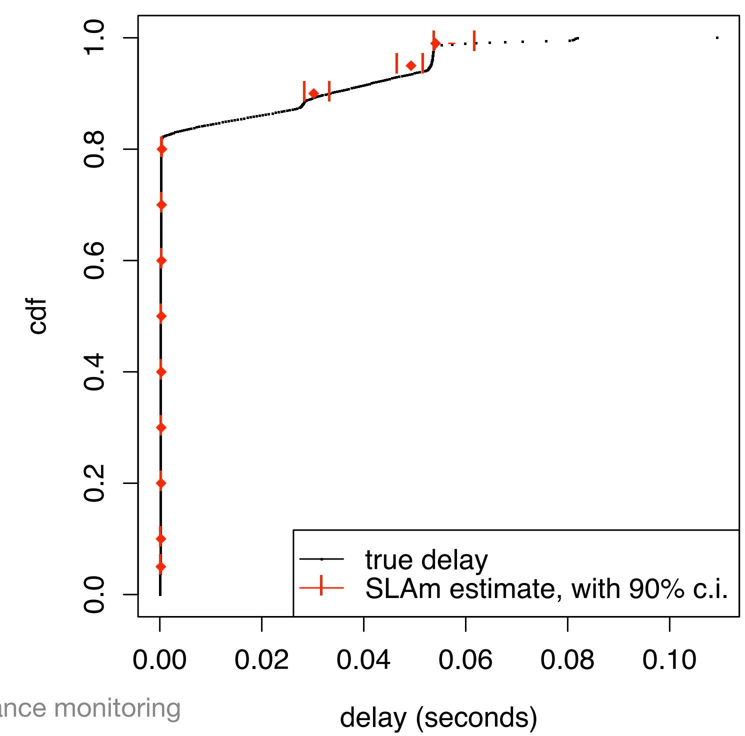

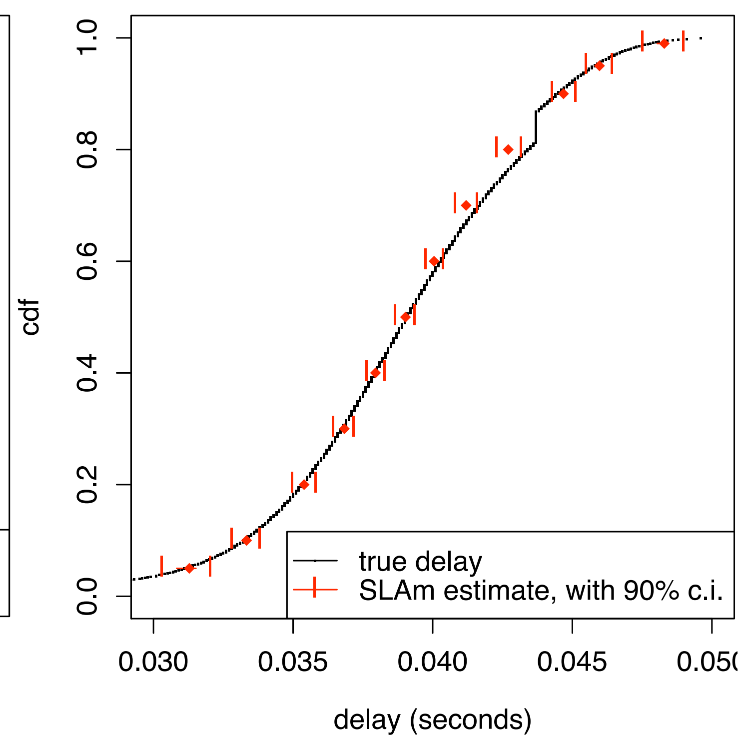

Results: Delay Distribution Inference

- Inferred distributions are close to the true ones
  - Discretization of 100 microseconds for convolution
- Results shown for UDP CBR traffic scenario (top) and self-similar traffic scenario (bottom)

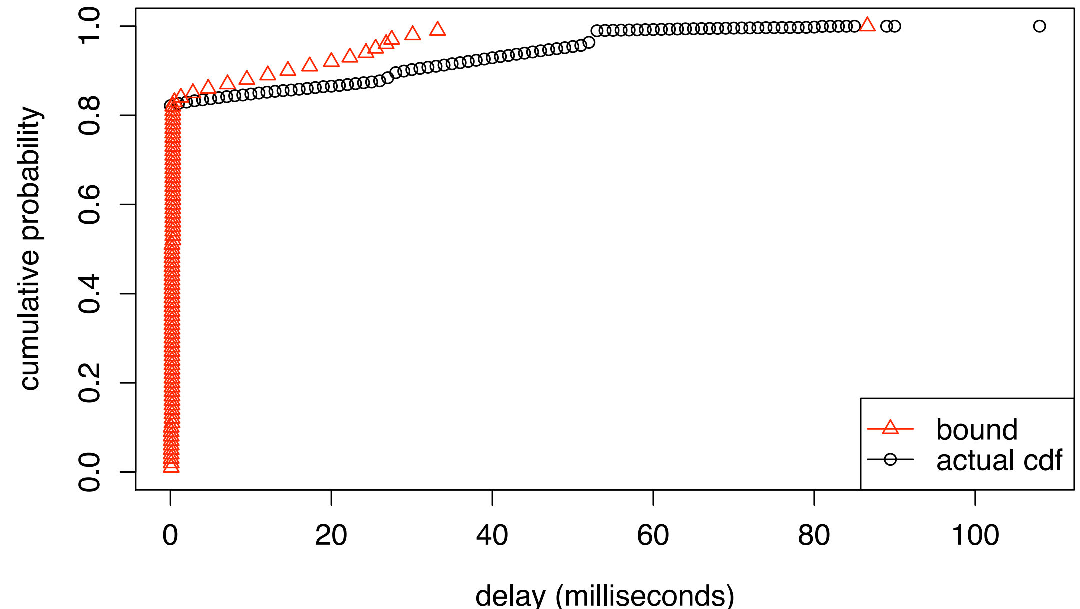

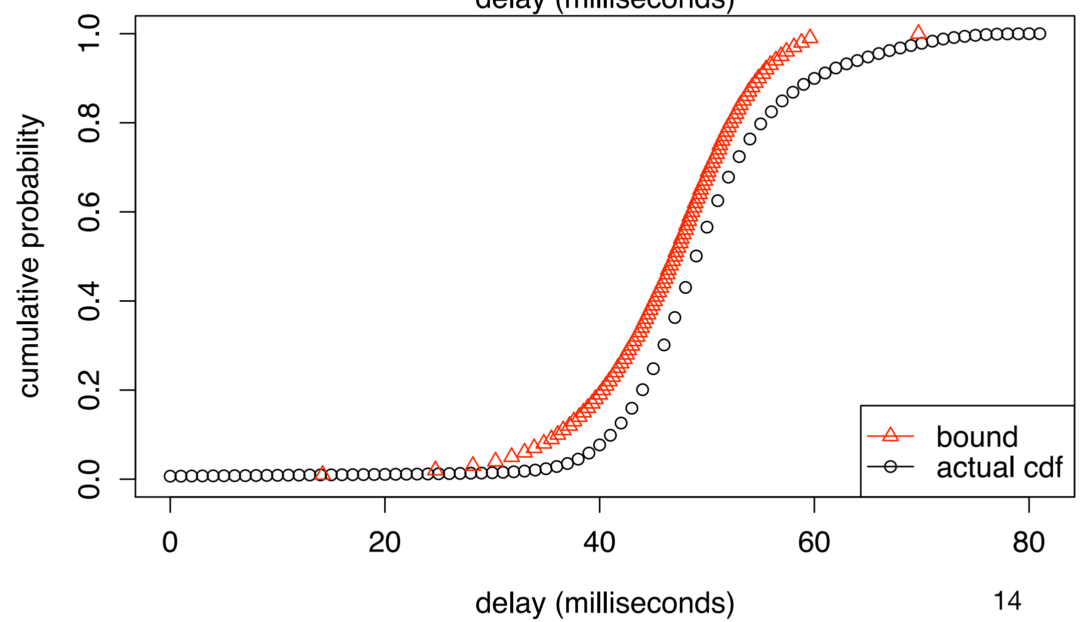

Results: Loss Rate

- Loss rate estimates are much more accurate than standard Poisson-based stream
  - Fast convergence to true loss rate (in paper)

Results for self-similar background traffic generated using Harpoon

loss rate comparison	SLAm		RFC 2680	
	true	estimate	true	estimate
dumbbell (60%)	0.0008	0.0007	0.0017	0
dumbbell (75%)	0.0049	0.0050	0.0055	0
star: route 1	0.0170	0.0205	0.0289	0.0058
star: route 2	0.0008	0.0006	0.0069	0.0000
star route 3	0.0192	0.0178	0.0219	0.0036
star: route 4	0.0005	0.0006	0.0002	0.0000

Results: Delay Variation

- SLAm DV matrix metric is more robust than RTP
  - More accurately tracks congested and turbulent conditions
  - Also robust in two-hop setting (in paper)