Accurate and Efficient SLA Compliance Monitoring

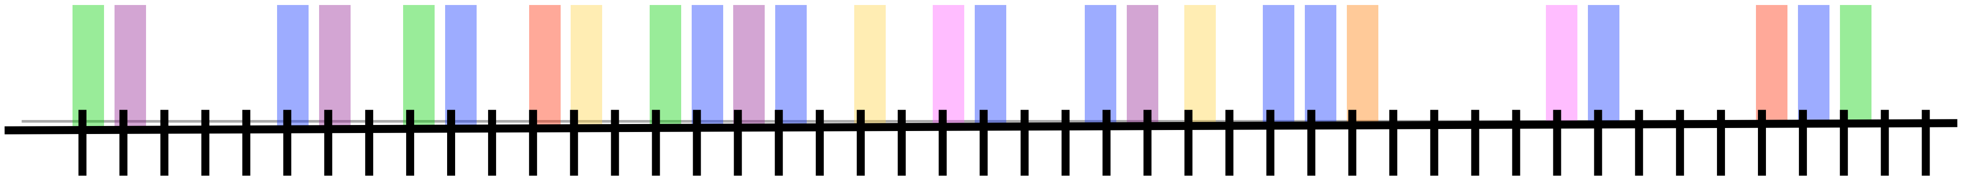

Joel Sommers

University of Wisconsin-Madison / Colgate University

Paul Barford

University of Wisconsin-Madison

Nick Duffield

AT&T Labs-Research

Amos Ron

University of Wisconsin-Madison

ACM SIGCOMM
Kyoto, Japan
29 August 2007

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!

Service Level Agreements

- Performance guarantees made by providers to customers.
  - *E.g.*, customer buys VPN service, wants guarantee of good service.
  - Metrics: packet loss, delay, delay variation, network availability.
  - Specific to origin-destination sites (*e.g.*, delay between A and C versus B and C).
  - Different statistics used, *e.g.*, mean, 95th percentile, maximum.
  - Metrics typically averaged over long time periods.

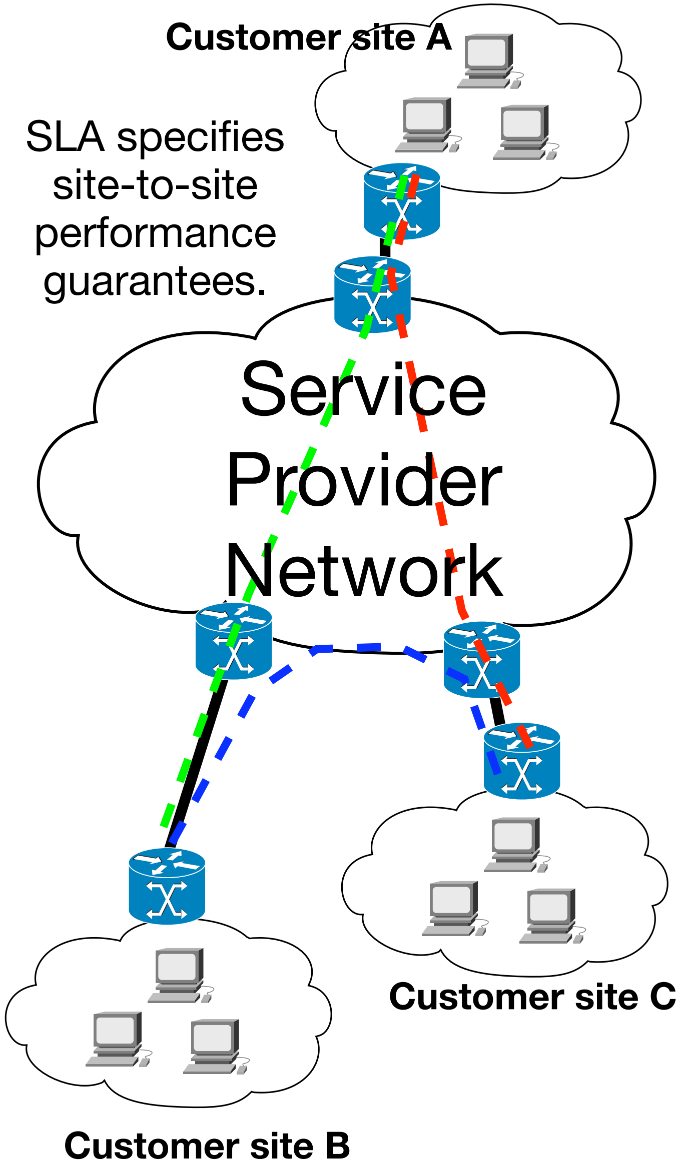

SLA Monitoring Challenges

- Overhead of simultaneous active measurement of multiple metrics is problematic.
- In-network characteristics are difficult to accurately measure with packet probes.
- Coordination and overhead of network-wide measurements.
- Data management.
    - Collection, processing, storing and archiving, coping with measurement errors, filtering outliers, ...

Approach

- Multi-objective probing: simultaneous measurement of multiple performance objectives.
  - Reduce overhead, simplify the measurement process.
- New and more accurate/robust active methodologies for measuring delay, loss, and delay variation
  - A new methodology for estimating mean end-to-end delay.
    - Based on Simpson's method for numerical integration.
  - A new methodology for estimating quantiles of the delay distribution.
    - No assumptions made about nature of the underlying distribution.

Approach (2)

- A new methodology for inferring an upper bound on the distribution of delay for an *unmeasured* path, given measurements for other, related paths in the network.
  - Extends algebraic approaches of prior work to *distributions*.
- A new heuristic for measurement of packet loss rate based on the badabing probe process [SBDR05].
  - Badabing originally designed to measure aspects of congestion episodes, not loss rate.
- A new methodology for more robust measurement of delay variation (jitter) on an end-to-end path.
  - A qualitative assessment of congestion, analogous to RTP.

Multi-objective probing

- Assume estimation algorithms operate in discrete time.
  - Probes may be scheduled to be sent at same time slot.
  - Tag probes according to the estimator module to which they apply.

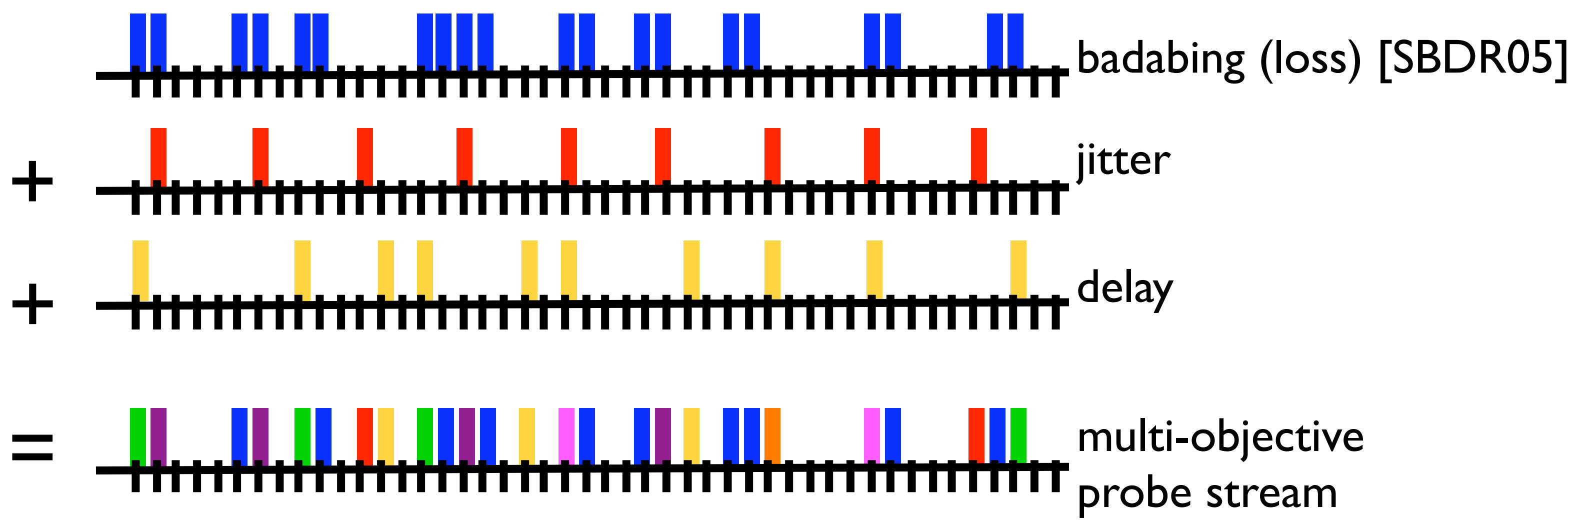

Implementation

- Discrete-time scheduler core.
  - Modular probe algorithms register with scheduler.
  - Module contains all logic to implement specific measurement algorithm.
  - Modules receive callbacks from scheduler, send probes through scheduler.

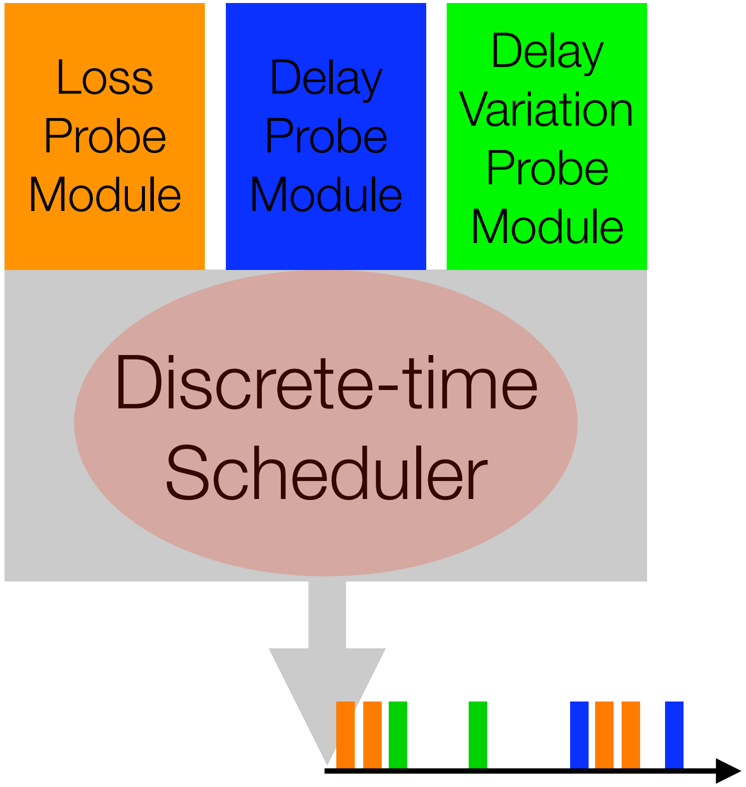

Probes tagged according to the estimation module.

Methodology: Mean Delay

- Model delay as a continuous function *f(t)*.
- Simpson's method for numerical integration is a natural approach for estimating the mean of *f(t)*.
  - *a, b* are the endpoints, and *c* is the midpoint of interval I_j.

$$\frac{1}{6}(f(a_j) + f(b_j) + 4f(c_j)) + e_j$$

- At time slot *i*, choose value *k* from geometric distribution with parameter *p*.
  - Send probes at time slot *i, i+(k+1), i+2(k+1)*.
  - Apply Simpson's method to measured probe delays.

Methodology: Delay Quantiles

- Estimate quantiles using delay samples from probes.
- Let $\{x_i : 1,\ldots,n\}$ be n samples drawn from distribution F.
  - Let Q_p denote the p^{th} quantile, the solution to $F(Q_p) = p$.
- $x_k \leq x$ is the event that at least k samples are less than or equal to x; $\Pr[x_k \leq Q_p] = G(n,p,k)$.
- Want: level $X^+(n,p,\epsilon)$ such that Q_p is guaranteed to exceed it with some small probability ϵ.
  - Use $K^+(n,p,\epsilon)$, the $1\text{-}\epsilon^{th}$ quantile of the binomial $B_{n,p}$ distribution.
  - Similar formulation for the lower bound $K^-(n,p, \epsilon)$.
  - Bounds can be calculated exactly using binomial distribution.

Methodology: Distribution Inference

- Consider scalar additive metrics (*e.g.*, delay, log transmission probability)
- Given a subset of performance measures across intersecting paths, is it possible to infer the whole set of measures?
    - Chen *et al*. (SIGCOMM 2004) and Chua *et al*. (INFOCOM 2005) examined problem for *scalar* measures.
- What about inferring a distribution of performance measures from a subset?

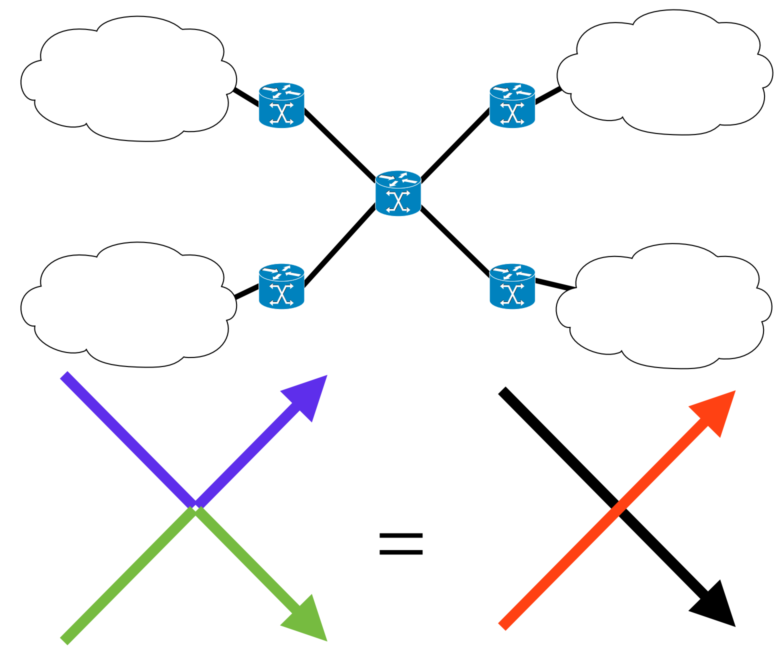

$y_1 + y_4 = y_2 + y_3$

(*e.g.*, y_i's are measured mean delays.)

Methodology: Distribution Inference

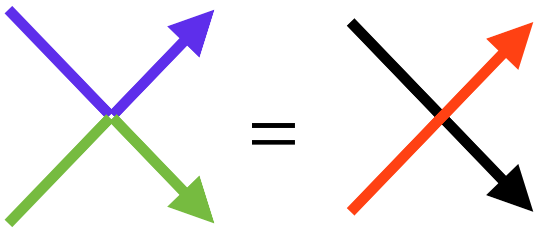

$Y_1 + Y_4 = Y_2 + Y_3$

(*e.g.*, Y_i's are distributions of one-way delays.)

- R is the set of routes forming routing matrix A.
- There is a minimal set of paths $S \subsetneq R$ s.t. every row of A can be expressed as a linear combination of S.
  - Partition S in S- and S+ based on sign of coefficient in the linear combination: $Y_1 = Y_2 + Y_3 - Y_4$; $S_1^+ = \{2,3\}$; $S_1^- = \{4\}$
    - Can formulate the convolution problem in terms of these partitions.
    - The distributions are discretized prior to convolution.
      - Our results provide a lower bound on the quantiles (upper bound on CDF).

Methodology: Loss Rate

- Start with badabing loss probe stream for measuring *frequency* and *duration* characteristics of loss episodes.
  - Probe pairs, sent according to a geometric distribution.
    - Each probe consists of three packets, sent back-to-back.
  - Heuristic: loss rate measured by badabing during a loss episode is related to what a typical TCP flow might measure.

$$\hat{L} = \hat{F}\hat{l}$$

Methodology: Delay Variation

- Consider a stream of probes of length *k*.
  - $s_{i,j}$ denotes difference in send time between probes *i,j*
  - $r_{i,j}$ denotes difference in receive time between probes *i,j*
- Construct a matrix *M* where element *i,j* contains the ratio $s_{i,j}/r_{i,j}$:
  - $s_{i,j}/r_{i,j} = 1$ if spacing does not change.
  - $s_{i,j}/r_{i,j} > 1$ if spacing increases.
  - $s_{i,j}/r_{i,j} < 1$ if spacing decreases.
  - $s_{i,j}/r_{i,j} = 0$ if either probe is lost.

Methodology: Delay Variation

- Compute eigenvalues of matrix *M*.
  - Results in vector *e* of eigenvalues, sorted large to small.
  - If all ratios are 1, largest eigenvalue is *k* (stream length).
    - Denote this “expected” vector of eigenvalues as *e’*.
  - Subtract *e’* from *e*, taking the L_1 norm of the resulting vector.

$$\sum_{i=1}^{k} |e_i - e'_i|.$$

- Result is called DV matrix metric.
  - A qualitative assessment of the amount of distortion from what we expect.

Experiments

- Created tool called SLAm (SLA monitor).
- Evaluated in controlled laboratory environment.
  - Two topologies: dumbbell and star.
- Compare SLAm with RFC standard probe streams at same bitrate.

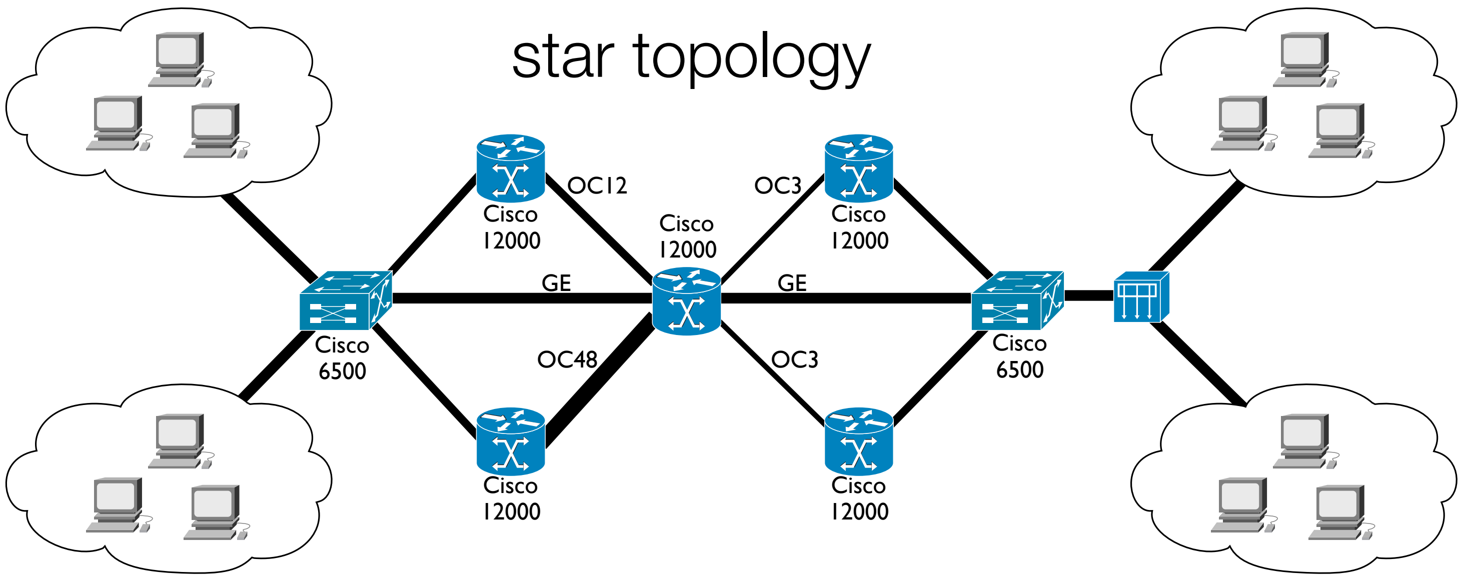

Results: Bandwidth Savings

- Three probe algorithms operating simultaneously.
  - 5 millisecond discrete time interval.
  - Loss probe: p_{loss} = 0.3, 600 byte packets.
  - Delay probe: p_{delay} = 0.048, 100 byte packets.
  - Delay variation periodic probe: 30 millisecond interval, 48 byte packets.
- Savings is parameter dependent, and can be big.

Loss	Delay	Delay Variation	Sum	SLAm	Savings
489 Kb/s	20 Kb/s	60 Kb/s	569 Kb/s	470 Kb/s	99 Kb/s (17%)

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).

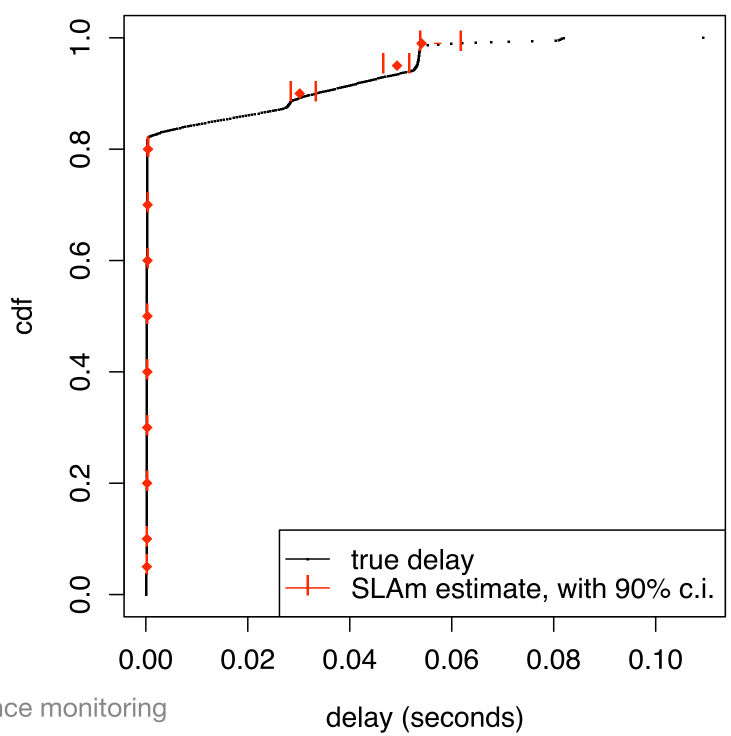

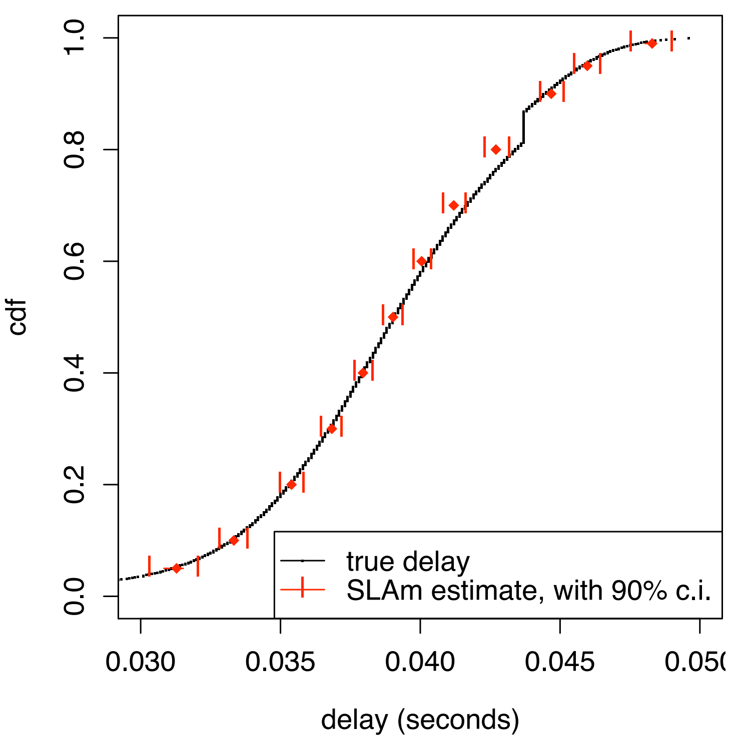

Results: Delay Distribution Inference

- 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).

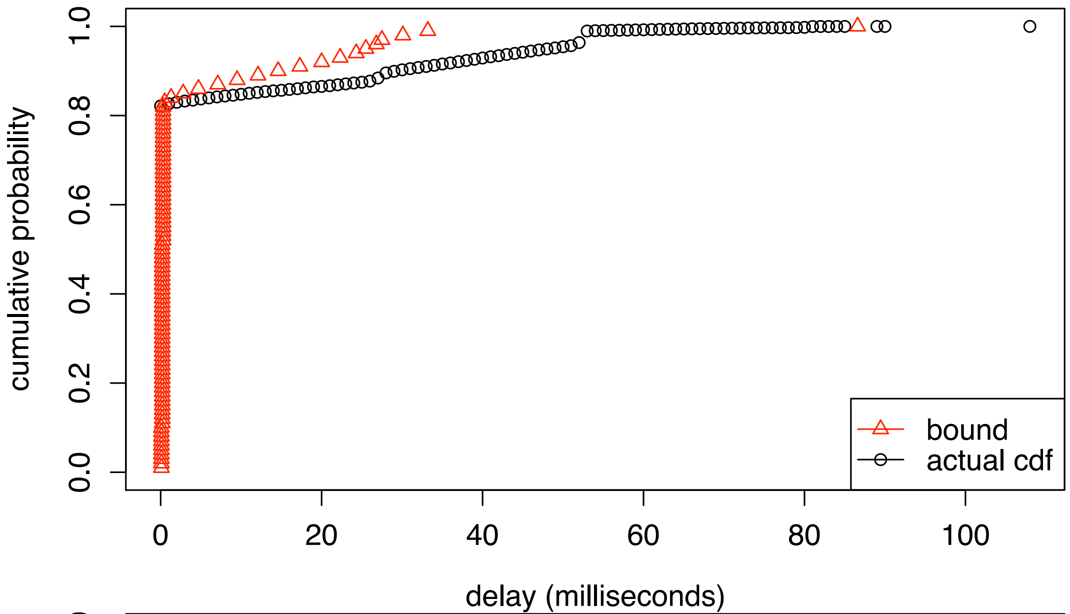

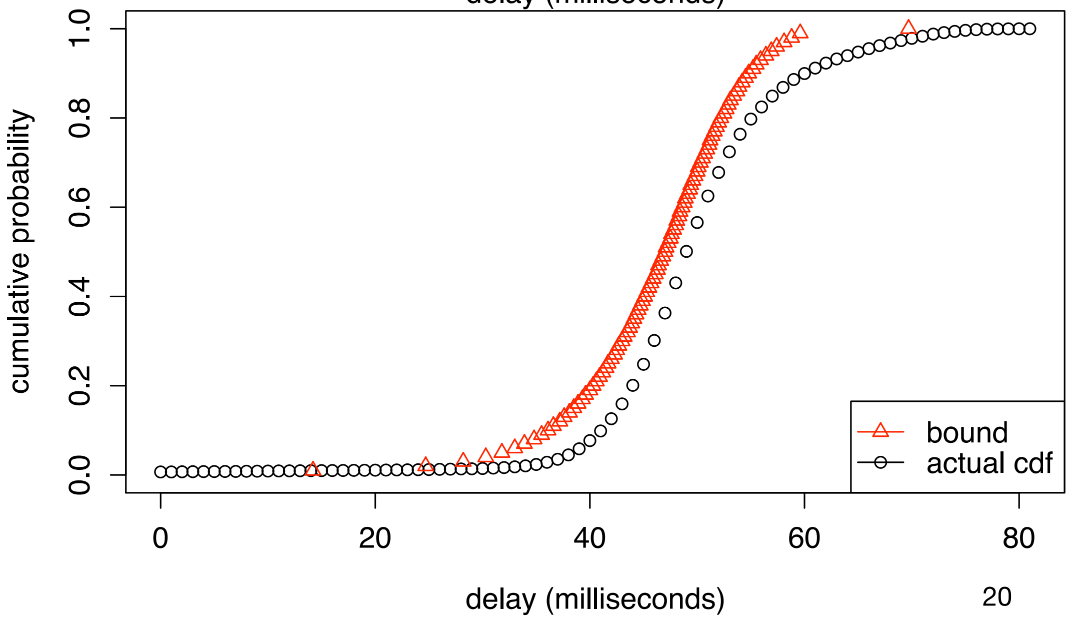

Results: Loss Rate

- 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).

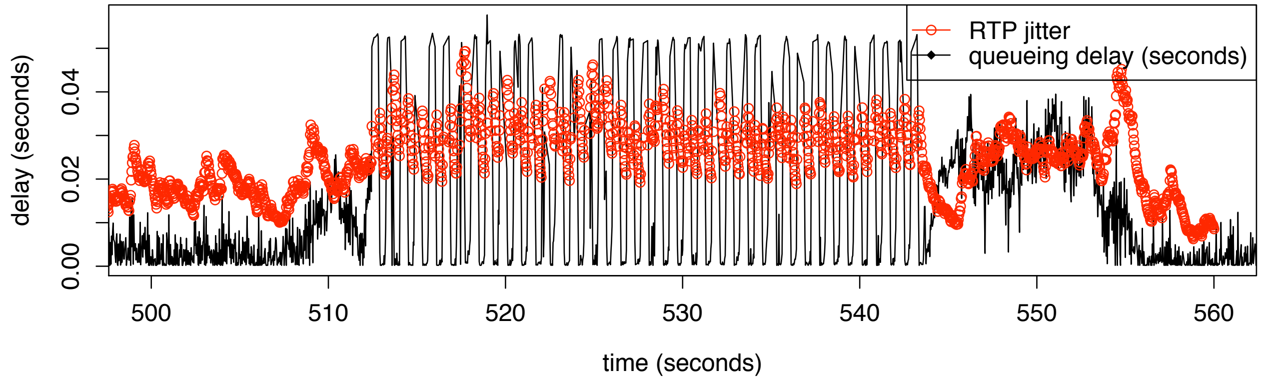

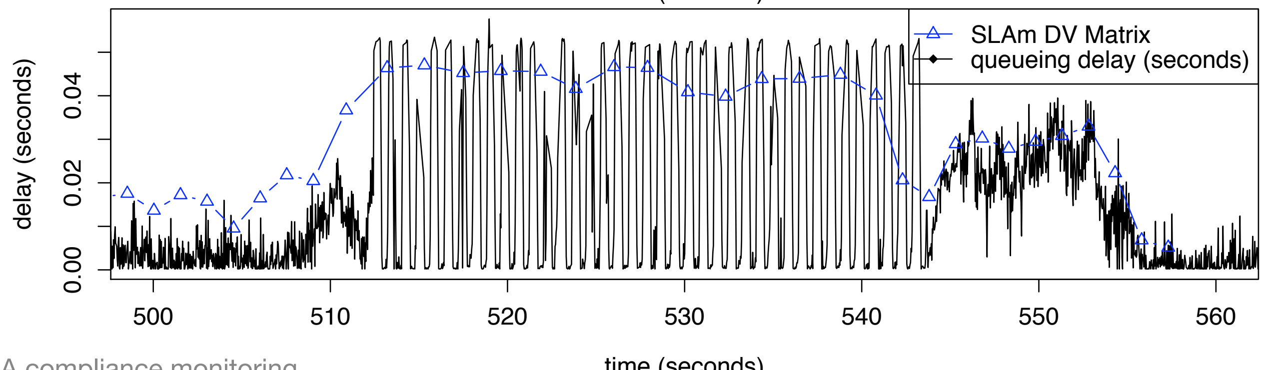

Summary

- A set of new methodologies for accurate, lightweight SLA compliance monitoring.
  - Multi-objective probing: reduces overhead.
  - Delay: accurate estimates of mean and quantiles; inferred distributions are close to true distributions.
  - Loss rate: accurate heuristic based on badabing probes.
  - Delay variation: robust qualitative estimate of congestion.
- Methodologies implemented in a tool called SLAm.
  - Laboratory tests with one- and two-hop topologies.
  - Source code will be released soon.

The end

- Ongoing and future work:
  - Probe stream coordination in the network-wide setting based on knowledge of topology.
  - How to optimize for accuracy given a daily (or hourly, etc.) probe budget?
  - SLA compliance monitoring does not require perfect accuracy; what appropriate tradeoffs be made between “good enough” accuracy and overhead?