Multi-Hop Probing Asymptotics In Available Bandwidth Estimation: Stochastic Analysis

Xiliang Liu, Kaliappa Ravindran
City University of New York
Dmitri Loguinov
Texas A&M University

Problem Statement

- Bandwidth estimation using first-order statistics of packet-train output dispersions
  - Assume an N-hop path probed by packet trains of length n

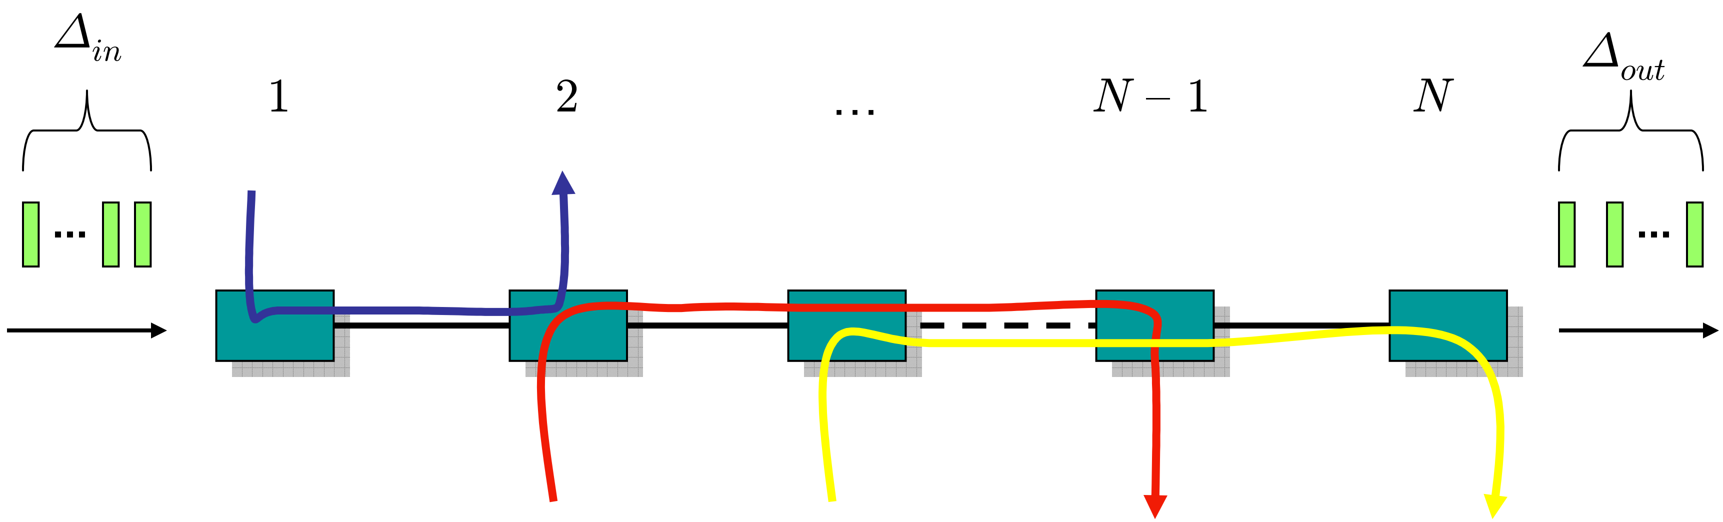

- **Goal**: derive the relationship between the statistical mean of Δ_{out} and Δ_{in} under arbitrary cross-traffic
  - We call this the probing response curve

Outline

<u>Related work and background</u>:

- Single-hop fluid curve
  - E.g., Melander (2001), Dovrolis (2001)
- Multi-hop fluid curve (one-hop persistent)
  - E.g., Dovrolis (2001)

<u>Our contribution</u>:

- Multi-hop fluid response curve
  - Arbitrary cross-traffic routing
- Multi-hop stochastic response curve
  - Packet-level model of cross-traffic
- Experimental verification
- Implications on existing techniques

Related Work

- Single-hop fluid setting:

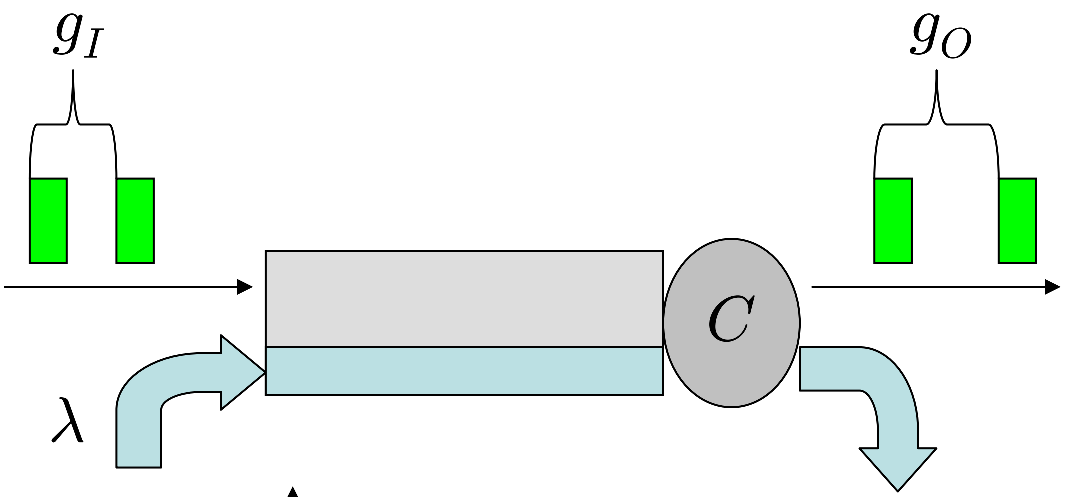

- Response curve:

$$g_O = \begin{cases} \frac{s}{C} + \frac{g_I \lambda}{C} & g_I \le \frac{s}{C-\lambda} \\ g_I & g_I \ge \frac{s}{C-\lambda} \end{cases}$$

$$= \max\left(g_I, \frac{s + g_I \lambda}{C}\right).$$

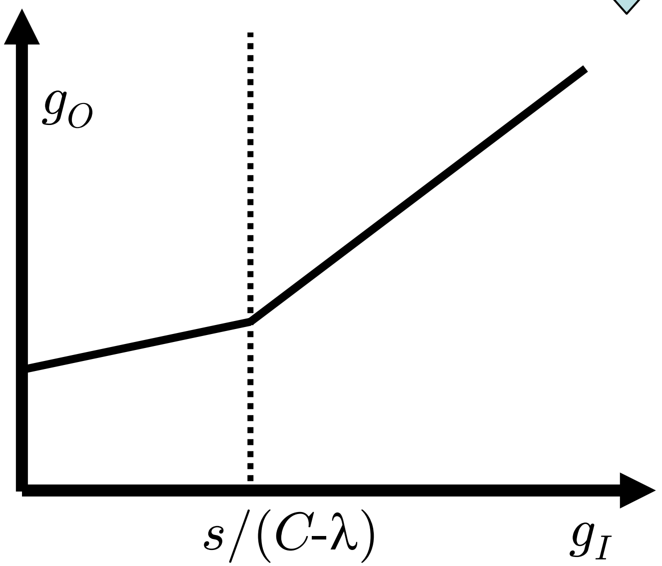

Related Work (cont'd)

Let $r_I = \frac{s}{g_I}$ and $r_O = \frac{s}{g_O}$, the fluid rate curve is

$$r_O = \begin{cases} r_I & r_I \leq C - \lambda \\ \frac{r_I}{r_I + \lambda} C & r_I \geq C - \lambda \end{cases} \quad (1)$$

A transformed version of rate curve depicts the relation between $\frac{r_I}{r_O}$ and r_I

$$\frac{r_I}{r_O} = \begin{cases} 1 & r_I \leq C - \lambda \\ \frac{r_I + \lambda}{C} & r_I \geq C - \lambda \end{cases} \quad (2)$$

Related Work (cont'd)

- Rate-response single-hop fluid curves:

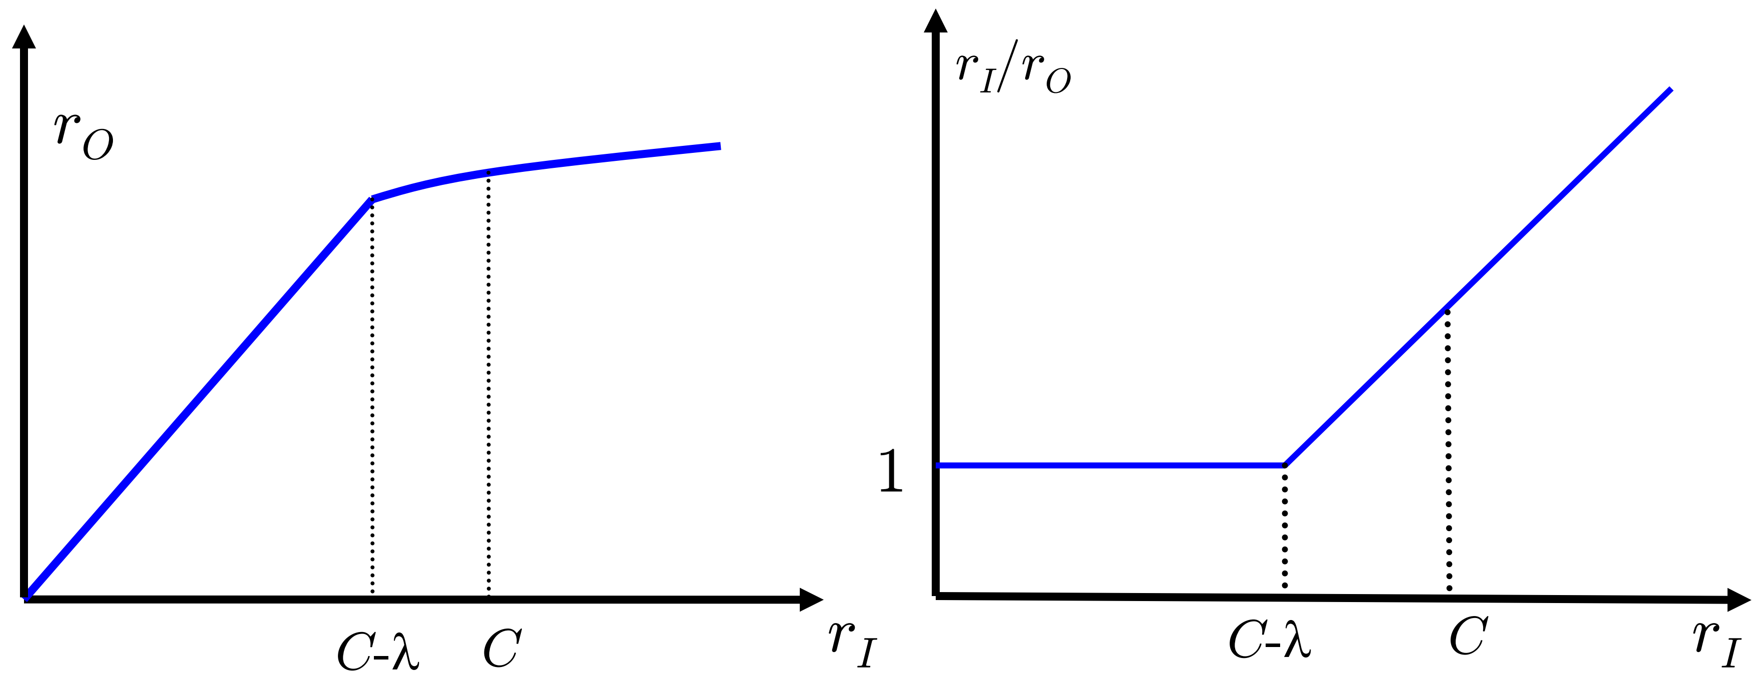

$$r_O = \min\left(r_I, \frac{r_I}{r_I+\lambda}C\right).$$

$$\frac{r_I}{r_O} = \max\left(1, \frac{r_I+\lambda}{C}\right).$$

Related Work (cont'd)

- Question: how do existing techniques relate to single-hop fluid curves?

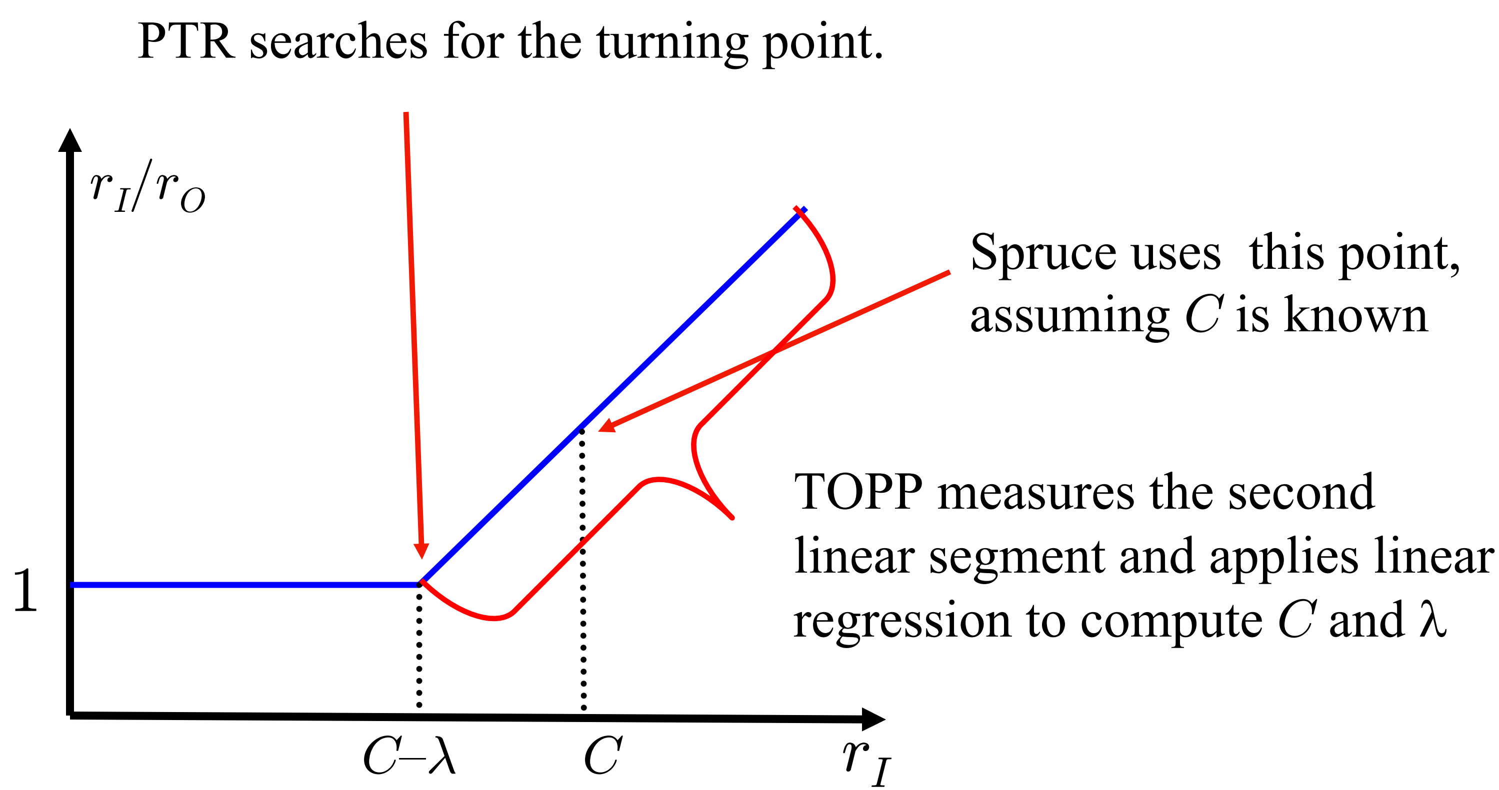

Related Work (cont'd)

- Previous multi-hop models
  - Analytical results are only available for fluid cross-traffic with one-hop persistent routing

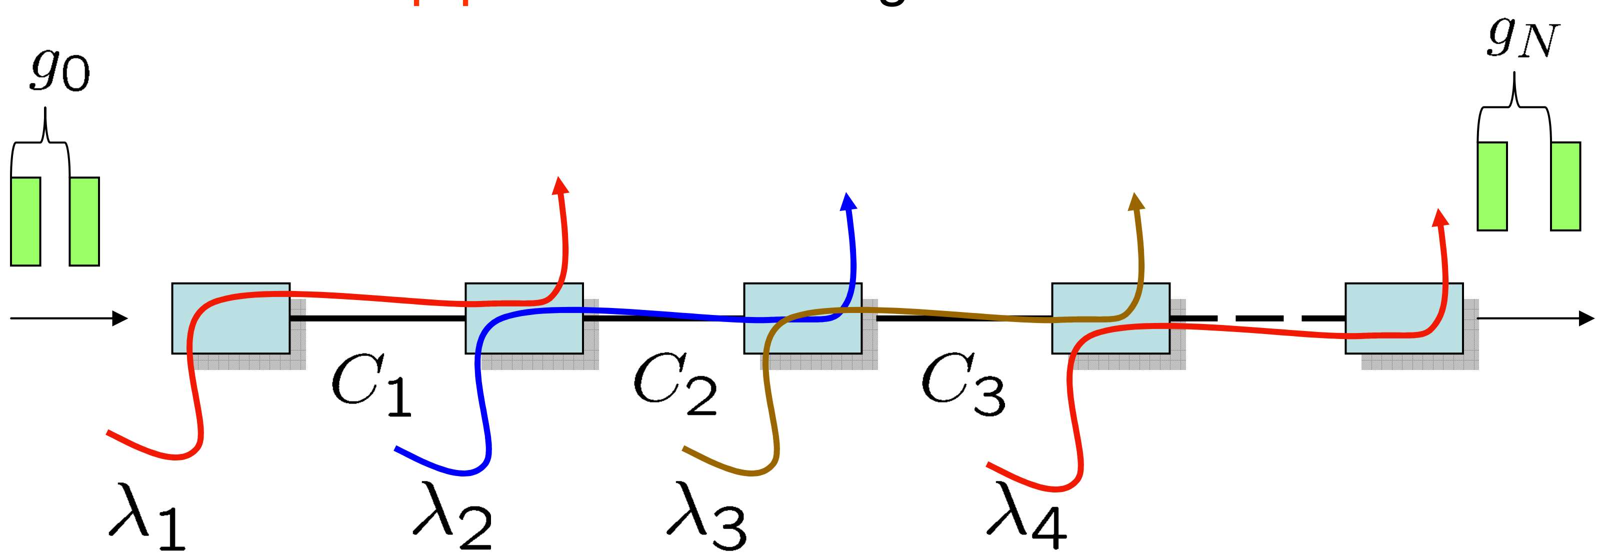

- Mathematically:

$$g_i = \max\left(g_{i-1}, \frac{s + \lambda_i g_{i-1}}{C_i}\right)$$

Outline

Related work and background:

- Single-hop fluid curve
  - E.g., Melander (2001), Dovrolis (2001)
- Multi-hop fluid curve (one-hop persistent)
  - E.g., Dovrolis (2001)

Our contribution:

- Multi-hop fluid response curve
  - Arbitrary cross-traffic routing
- Multi-hop stochastic response curve
  - Packet-level model of cross-traffic
- Experimental verification
- Implications on existing techniques

Multi-hop Fluid Response Curve

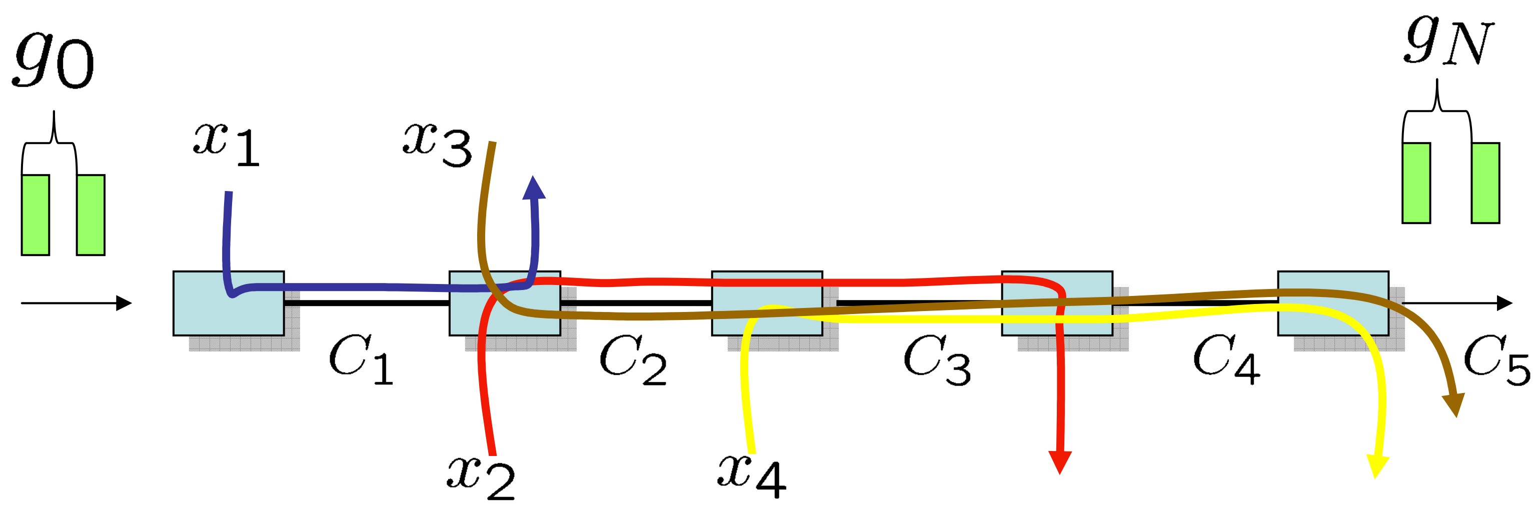

Notation 1: flow rate vector $\mathbf{x}$: $(x_1, x_2, \ldots x_M)$

Notation 2: flow aggregation $\Gamma_{k,i}$: A binary vector indicating all flows that enter the path at link k and traverse link i

Example: $\Gamma_{2,3} = (0, 1, 1, 0)^T$

The arrival rate of $\Gamma_{k,i}$ is $\mathbf{x}\Gamma_{k,i}$.

Multi-hop Fluid Response Curve (cont'd)

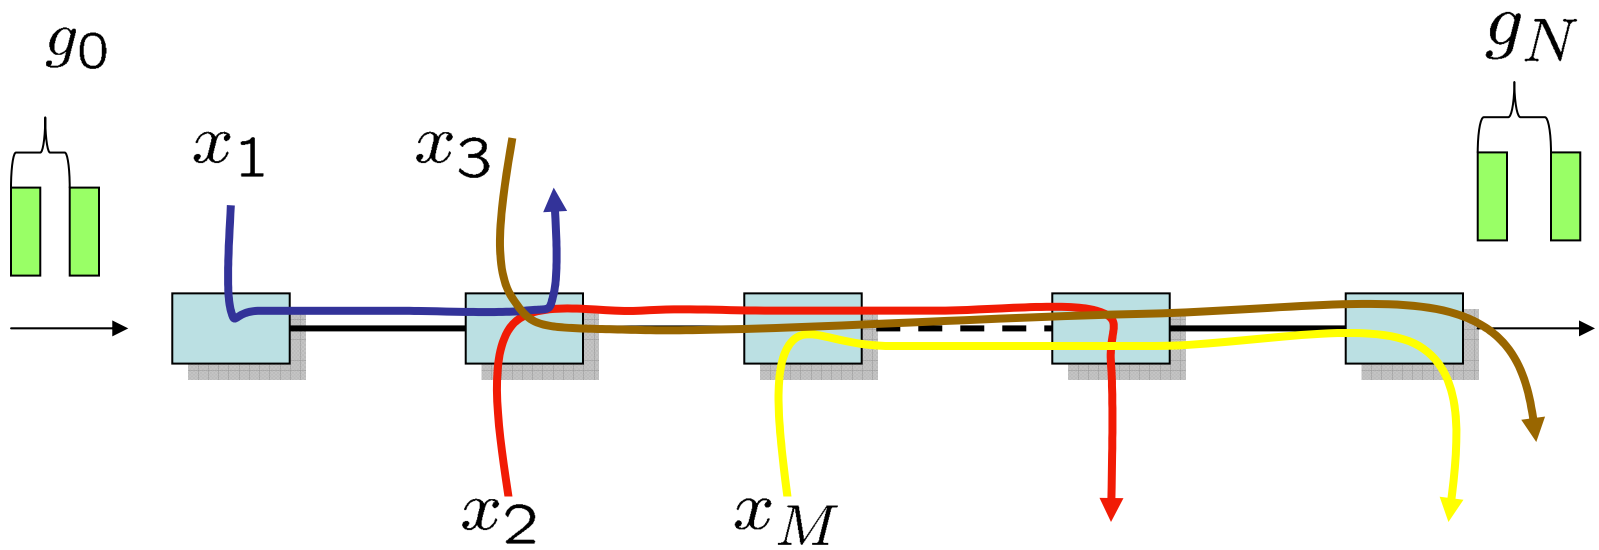

Main Result 1:

$$g_i = \max\left(g_{i-1}, \frac{s + \sum_{k=1}^{i} g_{k-1} \mathbf{x} \Gamma_{k,i}}{C_i}\right)$$

Multi-hop Fluid Response Curve (cont'd)

- Implications of this result
  - One-hop persistent curve is the upper bound
  - Single-hop curve is the lower bound

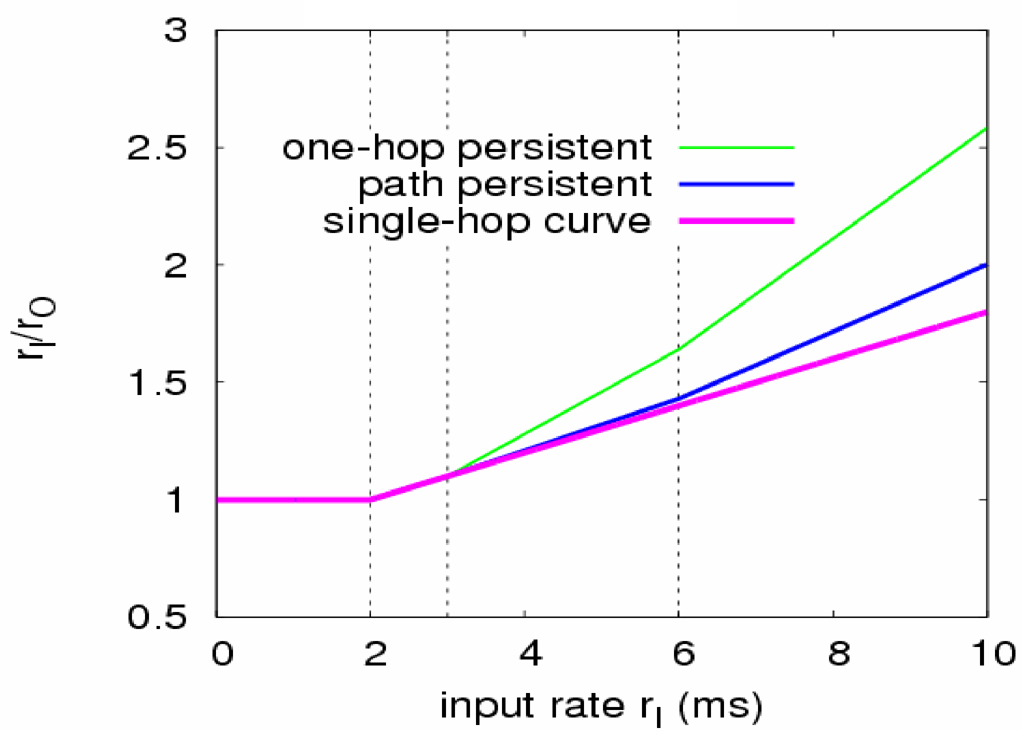

Outline

Related work and background:

- Single-hop fluid curve
  - E.g., Melander (2001), Dovrolis (2001)
- Multi-hop fluid curve (one-hop persistent)
  - E.g., Dovrolis (2001)

Our contribution:

- Multi-hop fluid response curve
  - Arbitrary cross-traffic routing
- Multi-hop stochastic response curve
  - Packet-level model of cross-traffic
- Experimental verification
- Implications on existing techniques

Multi-Hop Stochastic Response Curve

Assumption: The output dispersion R.V. G_N for different probings have idententical distributions.

Main Result 2:

$$
\begin{aligned}
E[G_i] &= E[G_{i-1}] + \frac{E[R_i]}{n-1} \\
&= \frac{\sum_{k=1}^{i} \left(\mathbf{x} \Gamma_{k,i} E[G_{k-1}] \right) + s}{C_i} + \frac{E[\tilde{I}_i]}{n-1}.
\end{aligned}
$$

Multi-Hop Stochastic Response Curve (cont'd)

- Stochastic curve:

$$E[G_i] = E[G_{i-1}] + \frac{E[R_i]}{n-1}$$
$$= \frac{\sum_{k=1}^{i}\left(\mathbf{x}\Gamma_{k,i}E[G_{k-1}]\right) + s}{C_i} + \frac{E[\tilde{I}_i]}{n-1}.$$

- Fluid curve:

$$g_i = \max\left(g_{i-1}, \frac{\sum_{k=1}^{i}\mathbf{x}\Gamma_{k,i}g_{k-1} + s}{C_i}\right).$$

These terms do not have fluid counterparts.

Multi-Hop Stochastic Response Curve (cont'd)

The presence of the two additional terms causes the multi-hop gap response curve to positively deviate from its fluid counterpart.

The response deviation $\beta_N = E[G_N] - g_N$ can be recursively expressed as follows

$$\beta_i = \begin{cases} \beta_{i-1} + \frac{E[R_i]}{n-1} & g_i = g_{i-1} \\ \frac{1}{C_i}\sum_{k=1}^{i}\left(\mathbf{x}\Gamma_{k,i}\beta_{k-1}\right) + \frac{E[\tilde{I}_i]}{n-1} & g_i > g_{i-1} \end{cases}$$

Multi-Hop Stochastic Response Curve (cont'd)

- Impact of Packet-train Parameters

Under certain additional assumptions, we get the following two results in the paper:

Main Result 3: For any given input rate, as probing packet size s increases, G_N converges to g_N in the mean.

Main Result 4: For any given input rate, as packet-train length n increases, G_N converges to g_N in the mean-square sense.

Full Picture of the 3 Response Curves

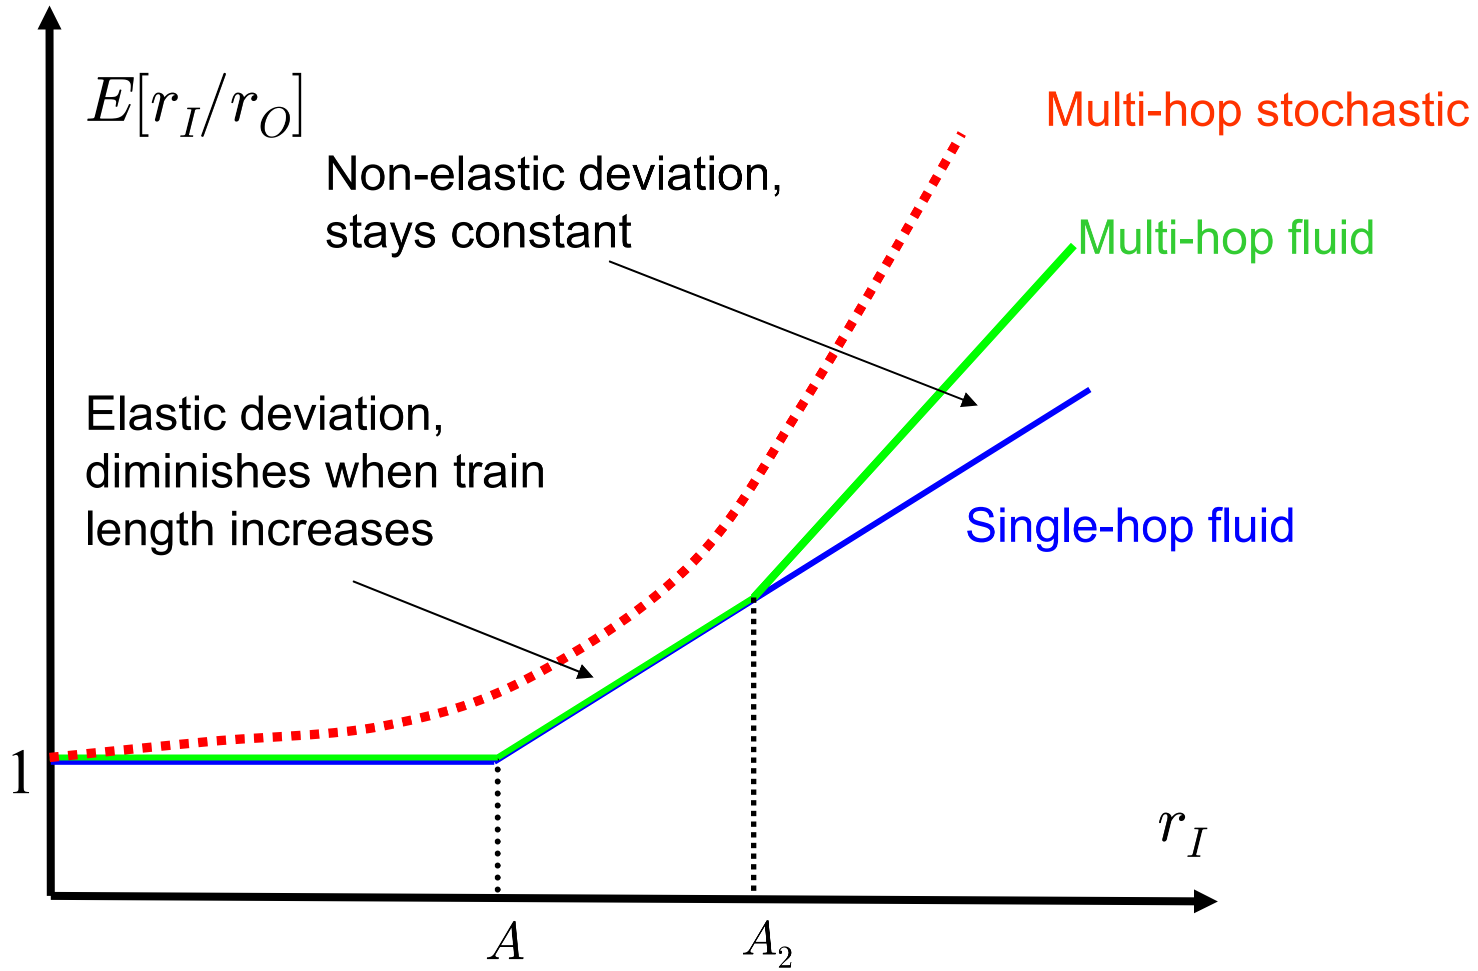

Outline

Related work and background:

- Single-hop fluid curve
  - E.g., Melander (2001), Dovrolis (2001)
- Multi-hop fluid curve (one-hop persistent)
  - E.g., Dovrolis (2001)

Our contribution:

- Multi-hop fluid response curve
  - Arbitrary cross-traffic routing
- Multi-hop stochastic response curve
  - Packet-level model of cross-traffic
- Experimental verification
- Implications on existing techniques

Experimental Verifications

- Emulab Testbed Settings

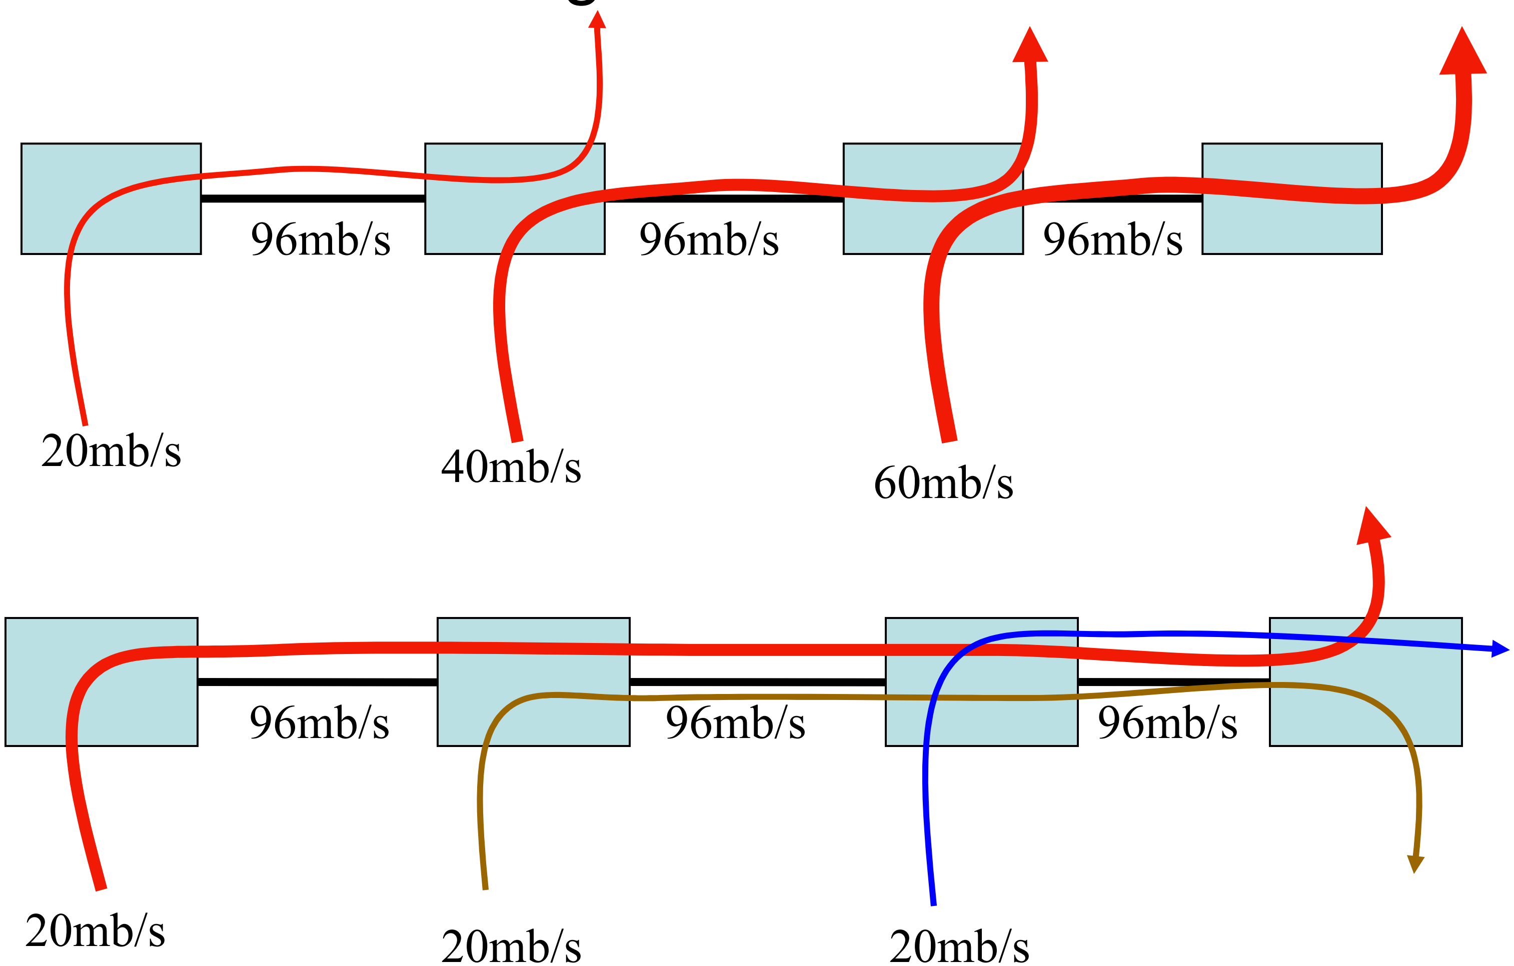

One-Hop Persistent Cross Traffic Routing

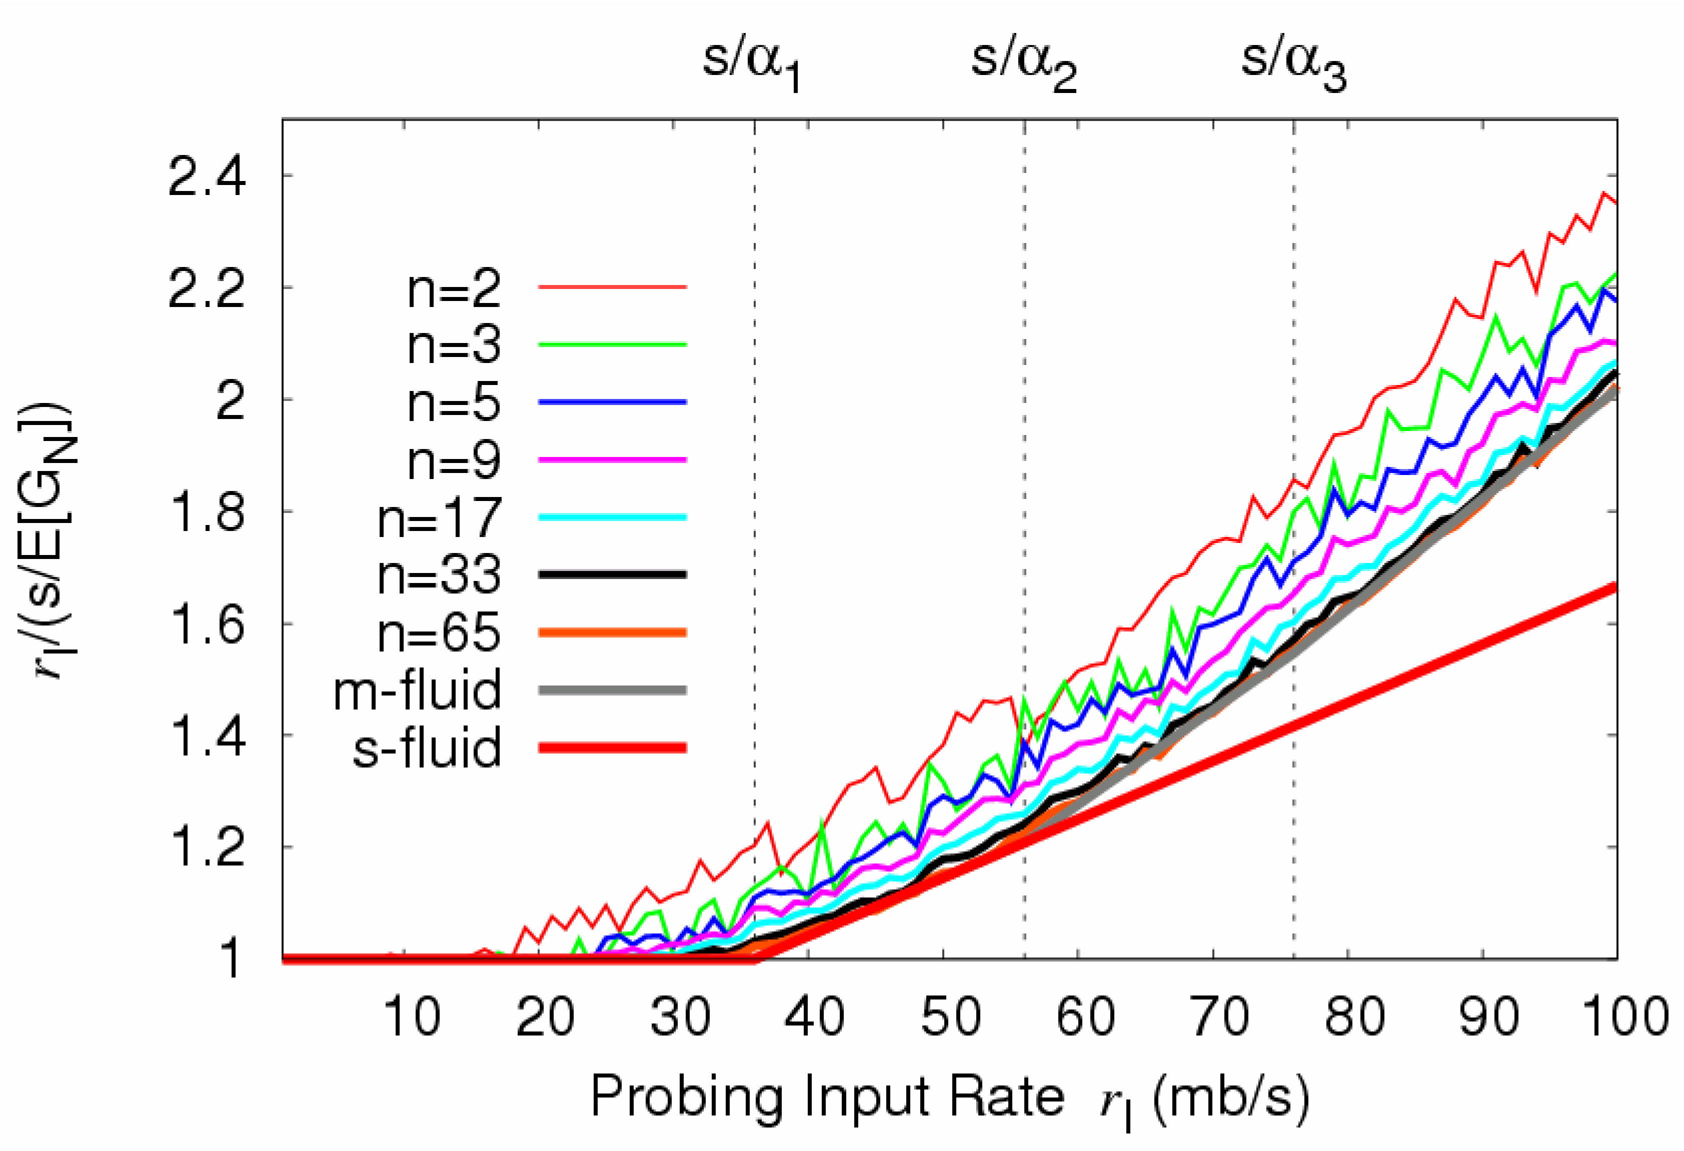

Path Persistent Routing

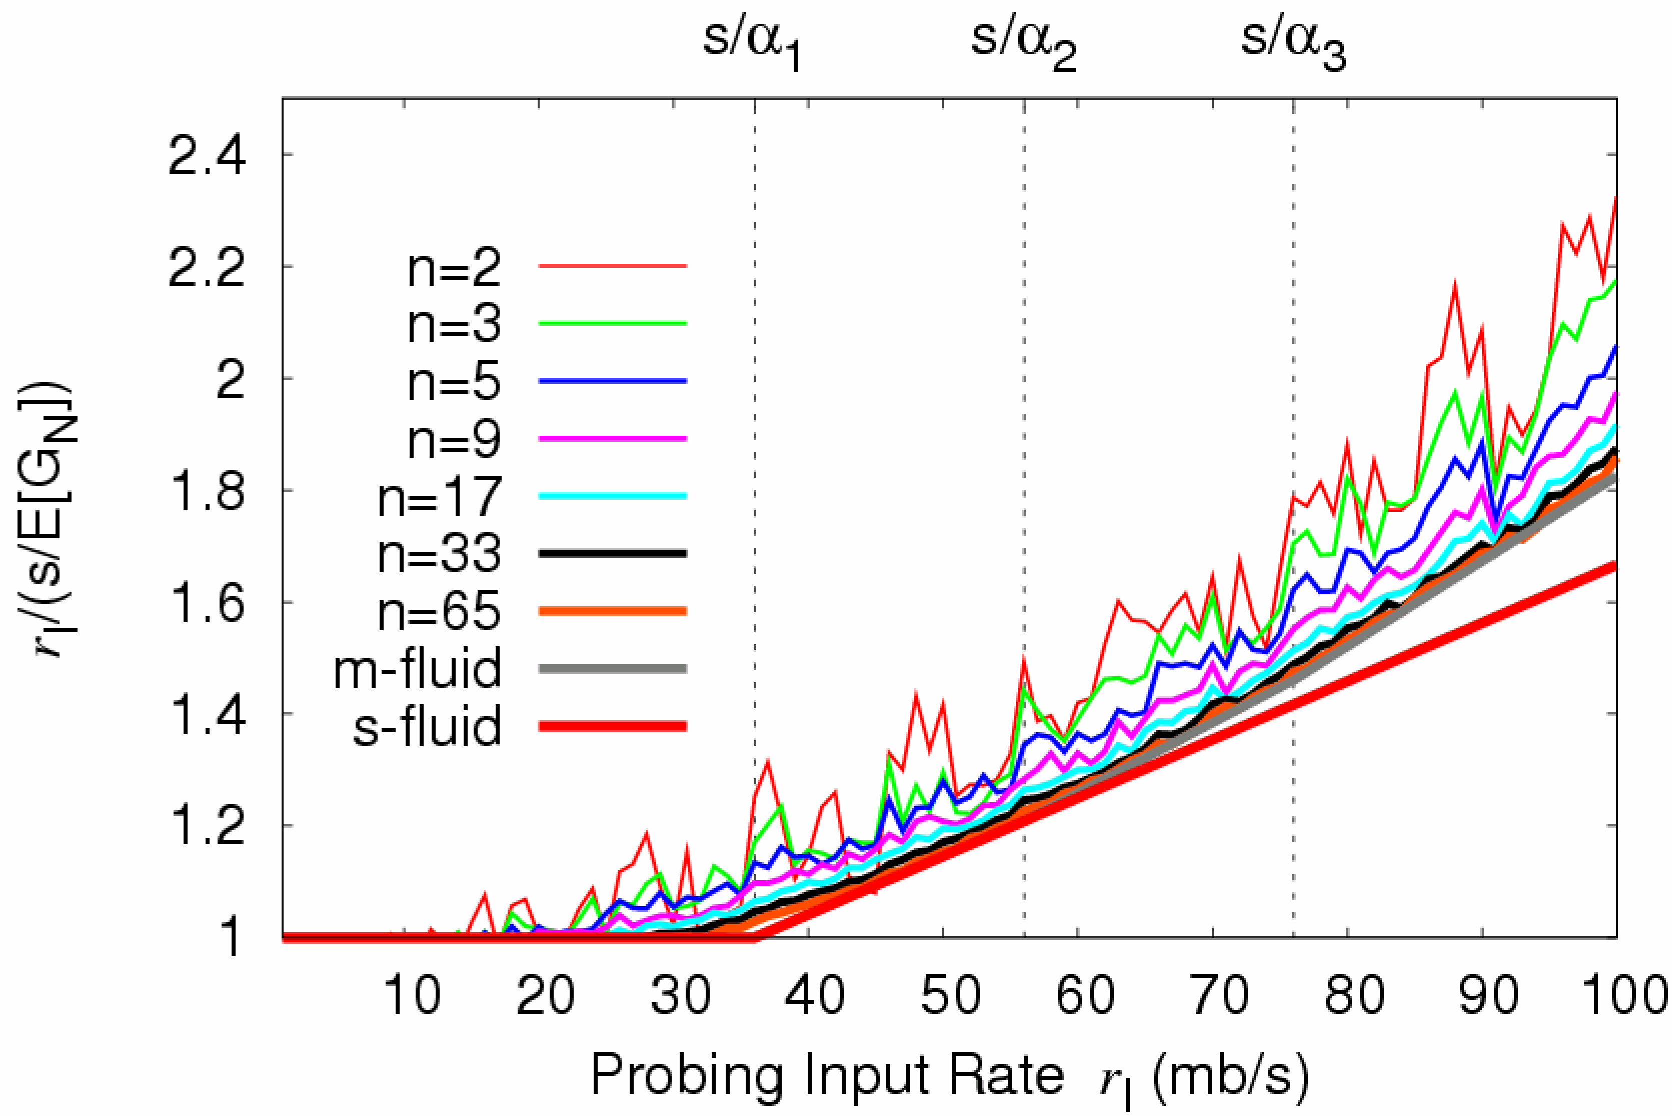

Real Internet Experiments

- We measure the rate response curves for more than 270 Internet paths over the RON testbed.
- Parameters:
  - Input rates: from 10 to 150 mb/s with step 5 mb/s
  - Packet-train length: from 33 to 129 packets
  - Packet-size: 1500 bytes
  - For each rate, we use 200 trains to estimate $E[G_N]$
- Experiment durations are 20-100 minutes

Lulea → CMU (1/16/2005)

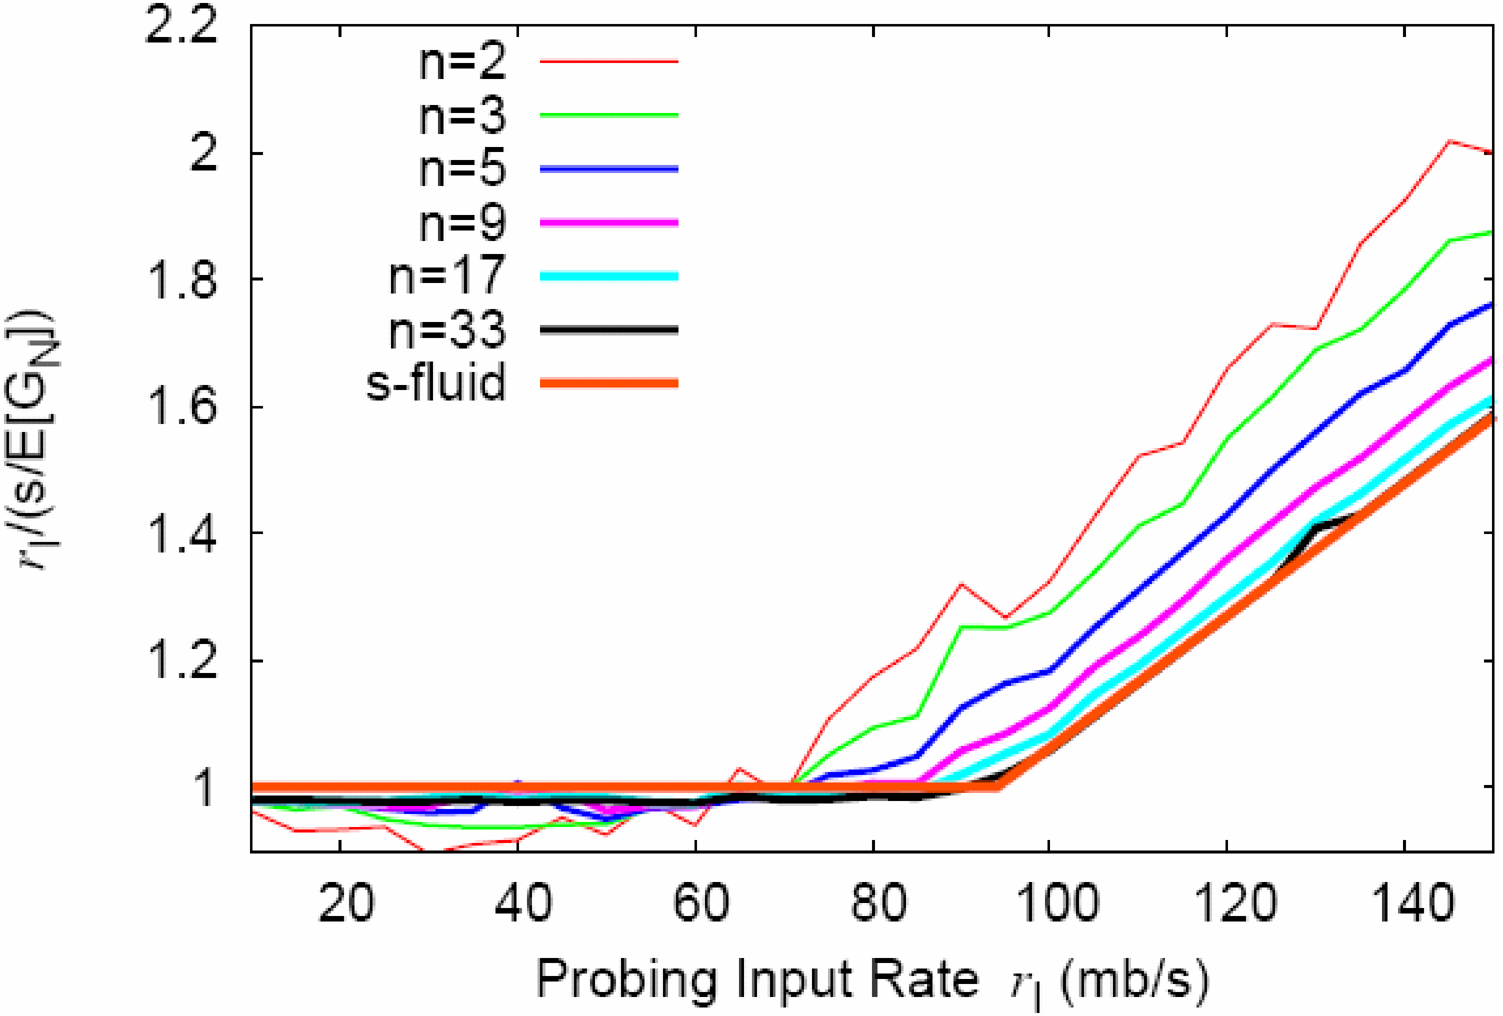

(a) lulea → CMU

Ana1-gblx → Cornell (4/29/2005)

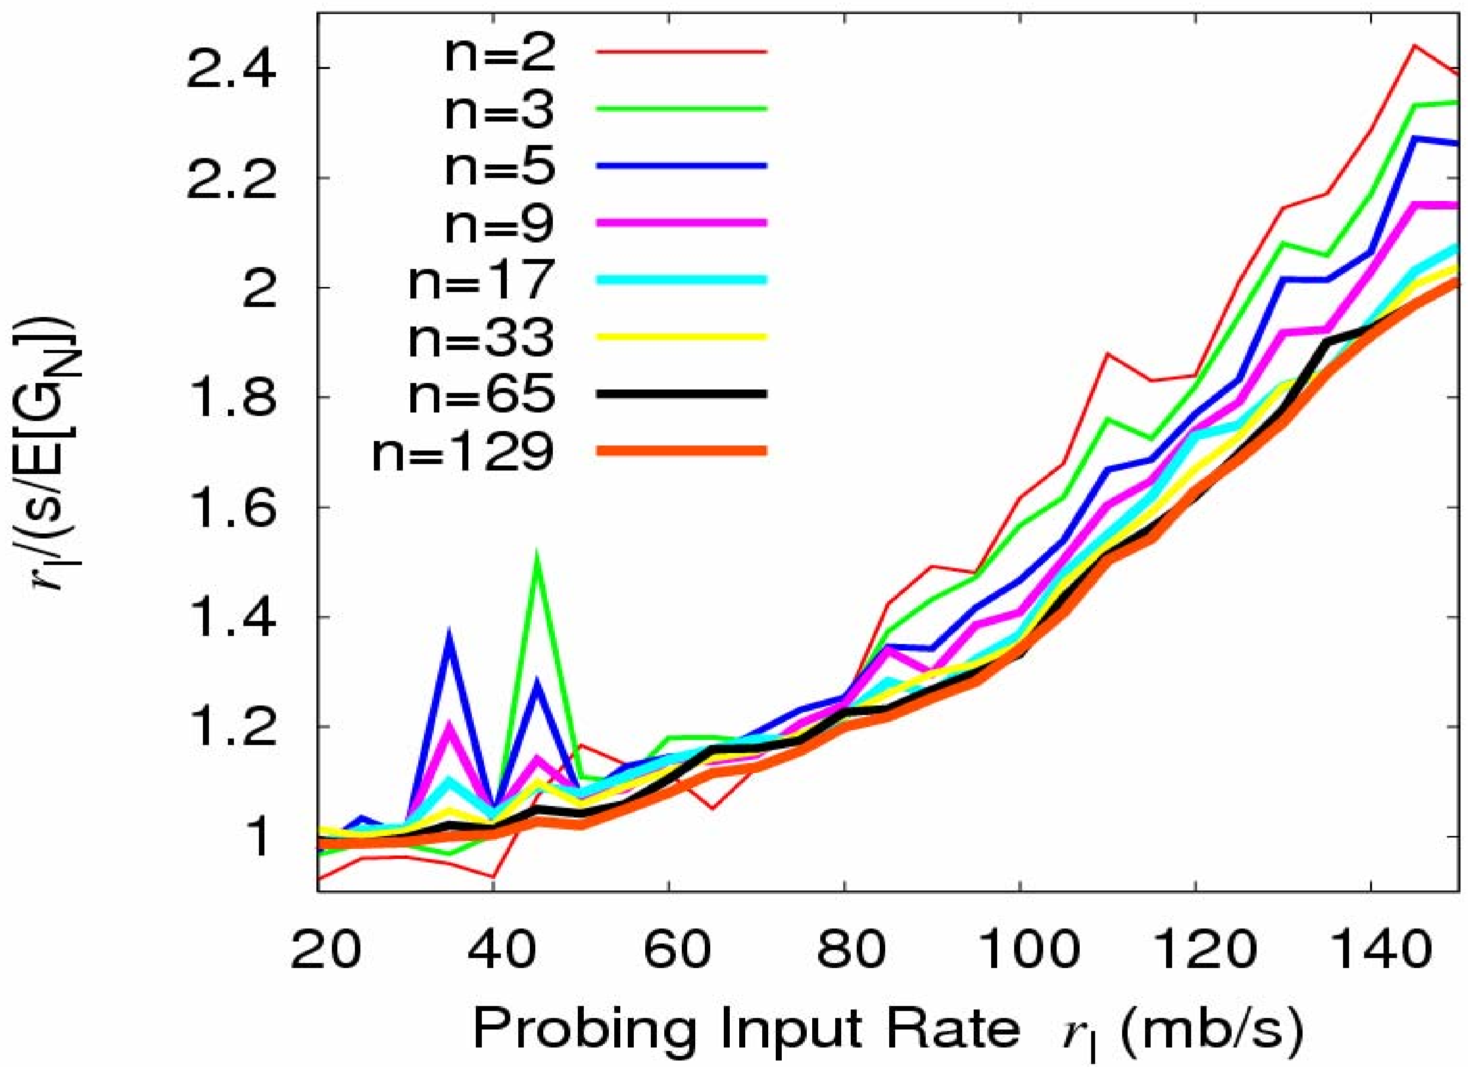

Outline

Related work and background:

- Single-hop fluid curve
  - E.g., Melander (2001), Dovrolis (2001)
- Multi-hop fluid curve (one-hop persistent)
  - E.g., Dovrolis (2001)

Our contribution:

- Multi-hop fluid response curve
  - Arbitrary cross-traffic routing
- Multi-hop stochastic response curve
  - Packet-level model of cross-traffic
- Experimental verification
- Implications on existing techniques

Implications on Existing Techniques

- TOPP uses packet-pairs to measure the stochastic response curve and implicitly assumes that it is the same as the fluid curve
  - To avoid bias, TOPP must use trains of sufficient length
- Pathload and PTR are related to searching for the turning point in the single-hop fluid response curve
  - Since they are using long trains, they are often immune to measurement bias

Implications on Existing Techniques (cont'd)

- Spruce measurement bias

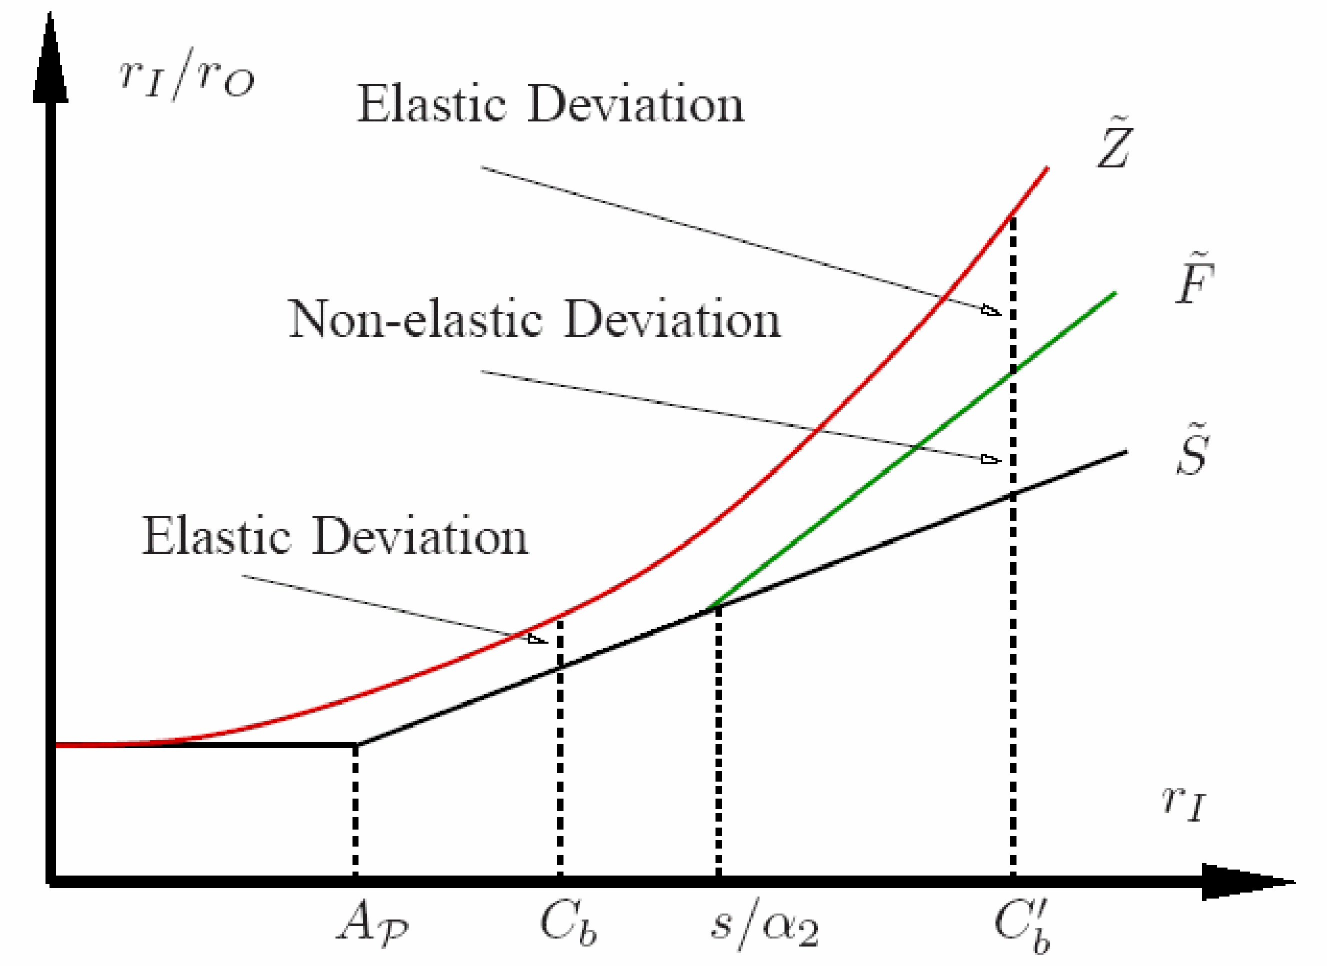

Implications on Existing Techniques (cont'd)

- Spruce measurement biases

Experiment	Elastic Bias	Non-Elastic Bias	Total bias	Real avail-bw	Spruce Measurement
Emulab-1	53.76	30.24	84	36	0
Emulab-2	26.88	12	38.88	36	0
Lulea-cmu	24	0	24	94	70

Conclusions

- We derived the multi-hop fluid response curve with arbitrary cross-traffic routing
- Also derived the multi-hop response curve using packet-based cross-traffic models and showed its convergence to its fluid counterpart when packet-train length increases

- Our results provided a stochastic justification of the existing techniques using long-trains
- Uncovered the sources of measurement biases for the techniques using short trains
- Leads to new techniques for measuring the tight link capacity (implementation in progress)