Bandwidth Estimation Using End-to-End Packet-Train Probing: Stochastic Foundation

Xiliang Liu

Joint work with Kaliappa Ravindran and Dmitri Loguinov

Department of Computer Science
City University of New York
New York, NY 10016

April 3, 2006

Outline

- Background
  - What is available bandwidth and why do we measure it ?
  - Single-hop fluid model and existing techniques
- Stochastic Foundation
  - Single-Hop case
  - Multi-Hop case
- Experimental Verification
- Implications

Background: Basics

- A network path: a number of packet-forwarding hops

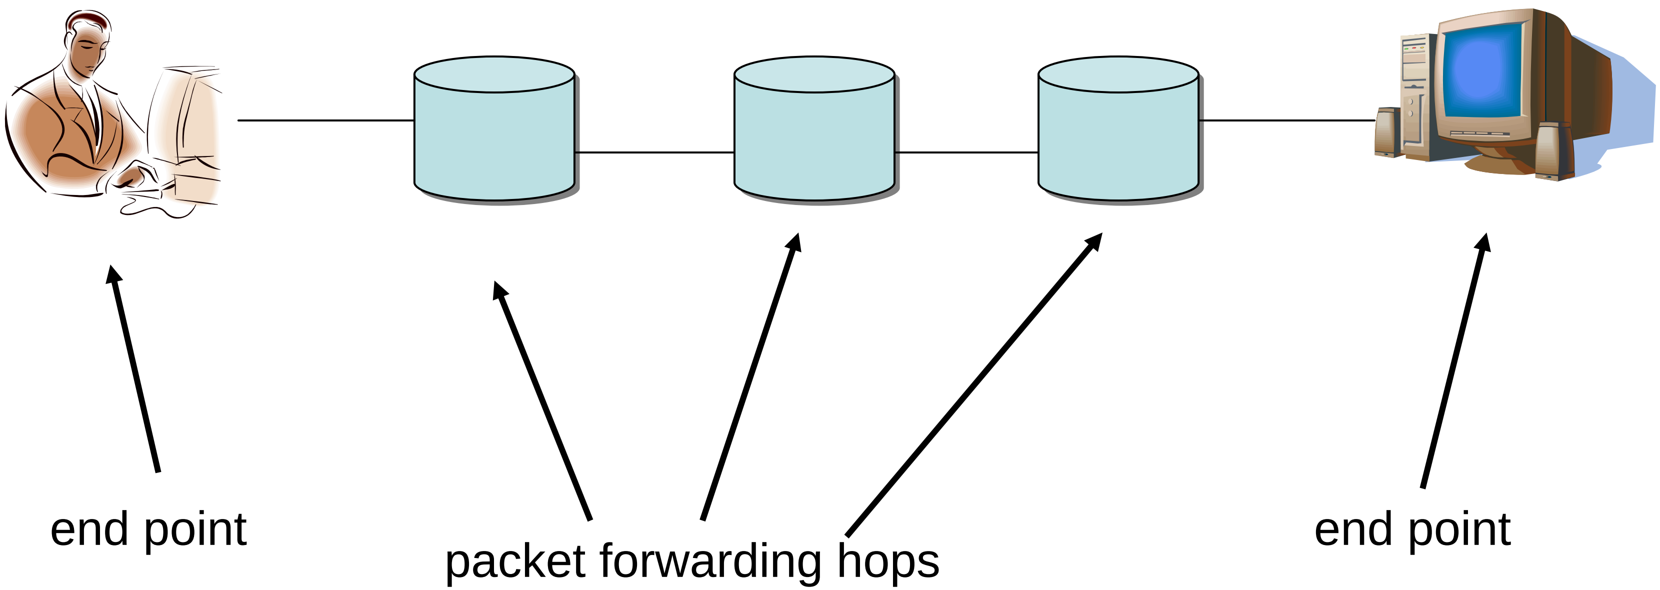

- End points: users of the path
  - Send packets to or receive packets from the path

Background: Basics

- A hop: a FIFO queue + a data transmission server

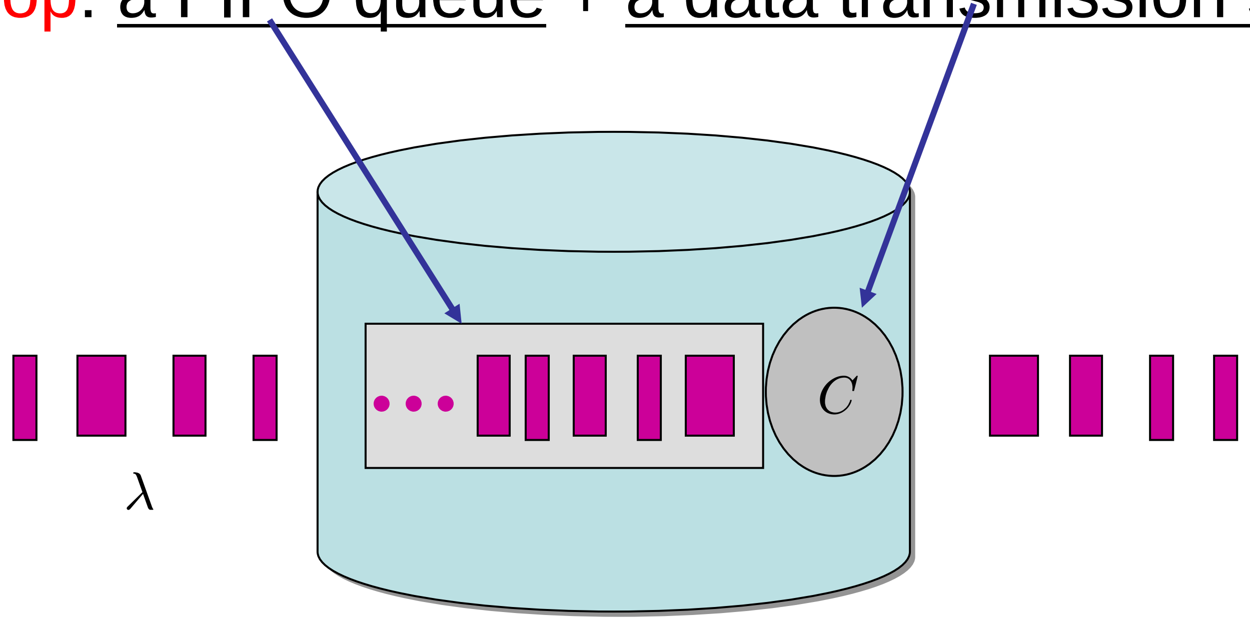

- Hop capacity C : transmission speed of the server in bits per second
- Cross-traffic rate λ : data arrived per time unit

Background: Definitions

- Available Bandwidth

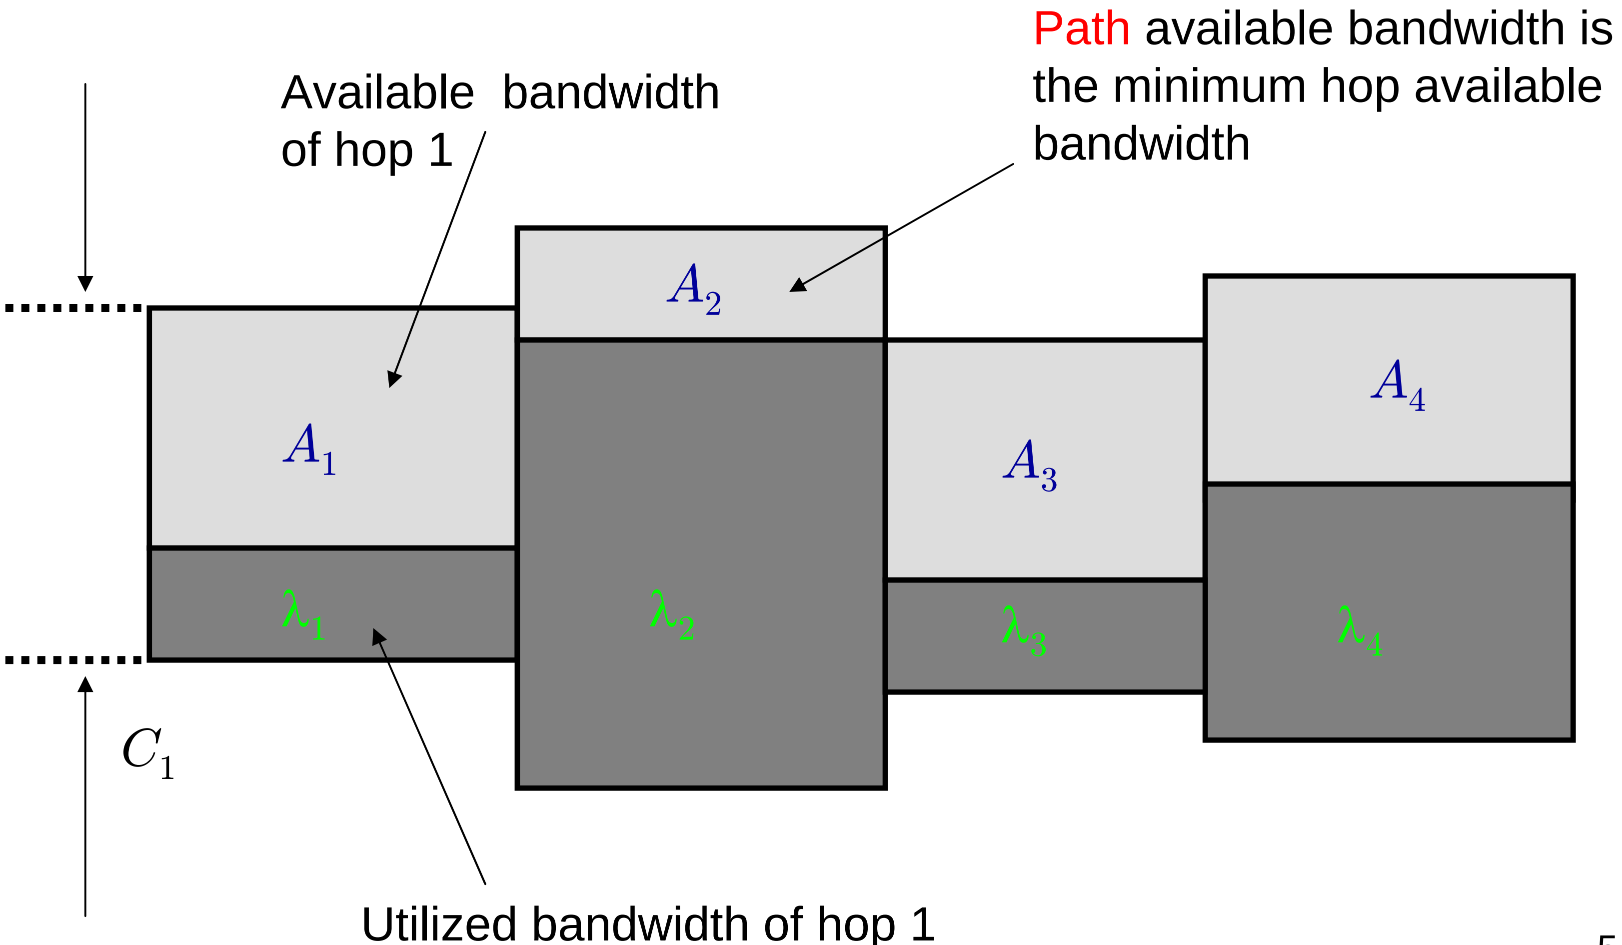

Background: Definitions

- Tight Link

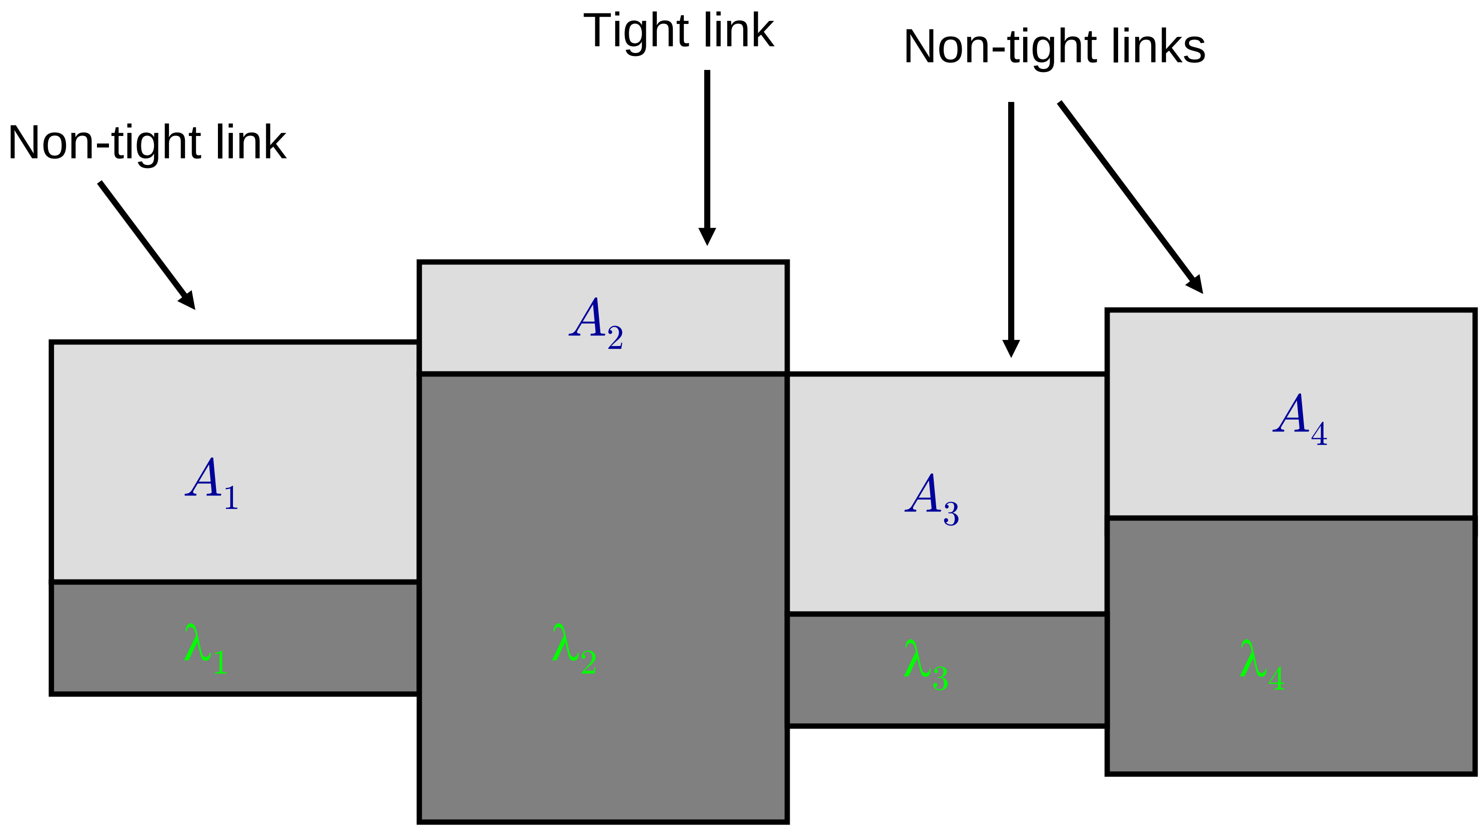

Background: Motivation

- Why measure available bandwidth?
  - Useful to a lot of applications
    - TCP ramp-up
    - Server selection
    - Overlay topology optimization
    - …
- Why measure from end points?
  - In the current Internet, end-users do not have access privileges of the data-forwarding hops

Background: Packet-Train Probing

- How to measure from endpoints?
  - Send probing packet-trains and infer bandwidth information from the input and output packet-train dispersions (i.e., the time gap between packets)

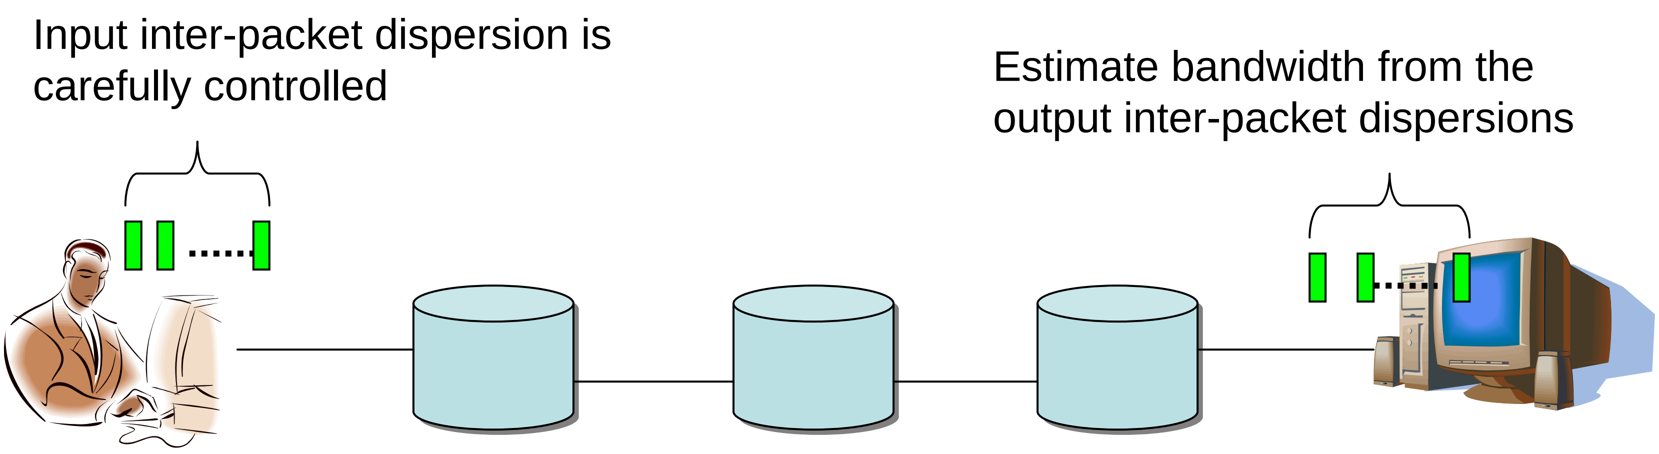

Background: More about Packet-Train

- Packet-train: a group of equally-sized packets

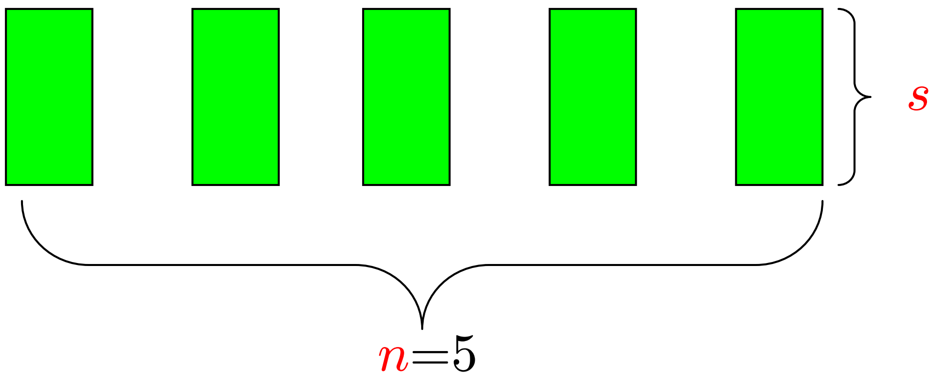

- Packet-train parameters
  - Size s and length n (not changeable by the path)

Background: More about Packet-Trains

- Packet-train: a group of equally-sized packets

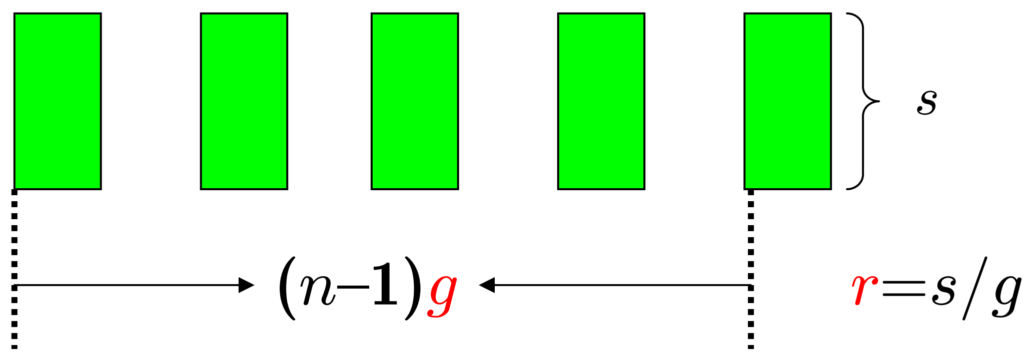

- Packet-train parameters
  - Size s and Length n (not changeable by the path)
- Signals carried by a packet-train
  - Dispersion g and rate r (changeable by the path)

Background: Single-Hop Fluid Model

- Most existing measurement techniques are designed based a single-hop fluid model

cross-traffic rate λ

- Fluid cross-traffic
  - Infinitely small packet size
  - Constant arrival rate λ at any time interval

- Path available bandwidth is $C–\lambda$ constantly

Background: Single-Hop Fluid Model

- Most existing measurement techniques are designed based a single-hop fluid model

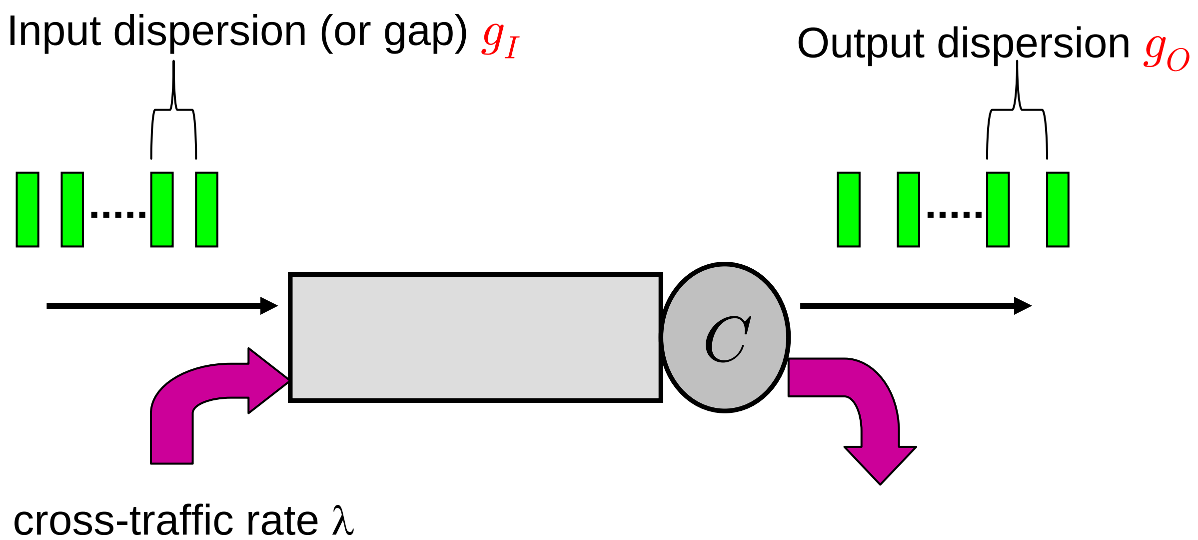

- Output: a response of the path to the input
  - Input-output relation is called the response curve of the path

Background: Single-Hop Fluid Curves

- Single-hop fluid gap response curve

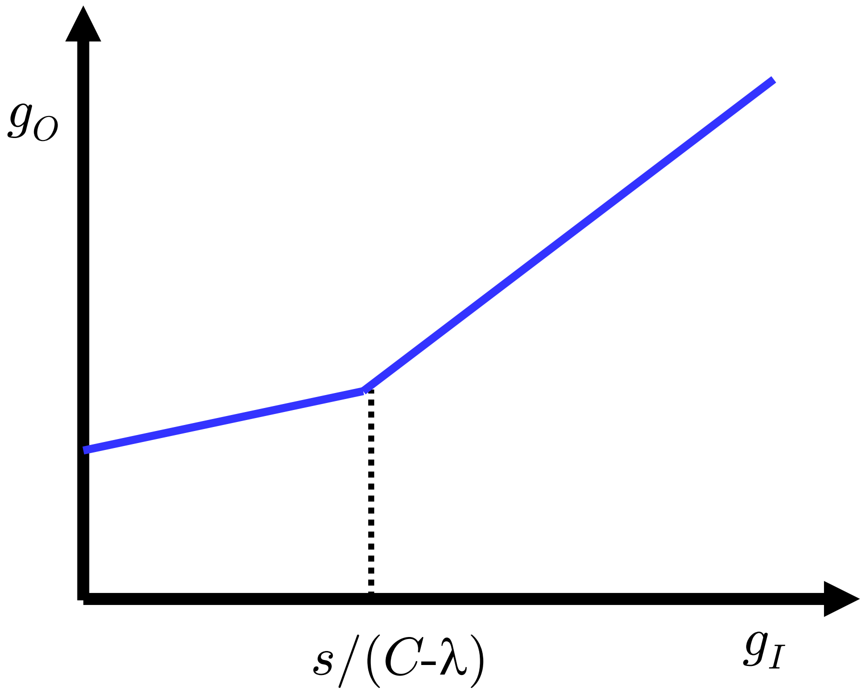

$$g_O = \begin{cases} \frac{g_I \lambda}{C} + \frac{s}{C} & g_I \le \frac{s}{C-\lambda} \\ g_I & g_I \ge \frac{s}{C-\lambda} \end{cases}$$

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

Background: Single-Hop Fluid Curves

- Explanation using a Hop Workload Graph
  - Workload: the amount of data waiting for transmission in the queuing system, measured in transmission time

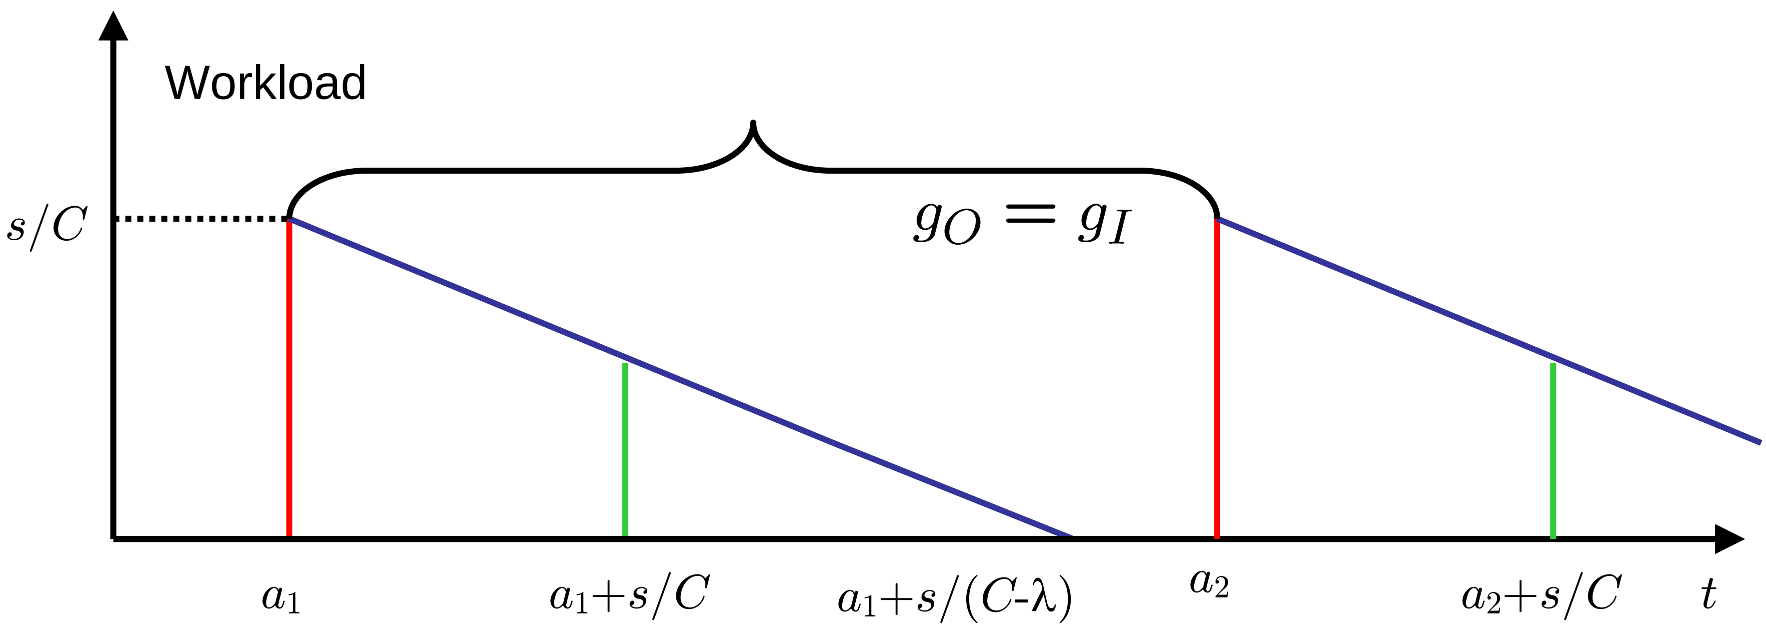

Background: Single-Hop Fluid Curves

- Explanation using a Hop Workload Graph
  - Workload: the amount of data waiting for transmission in the queuing system, measured in transmission time

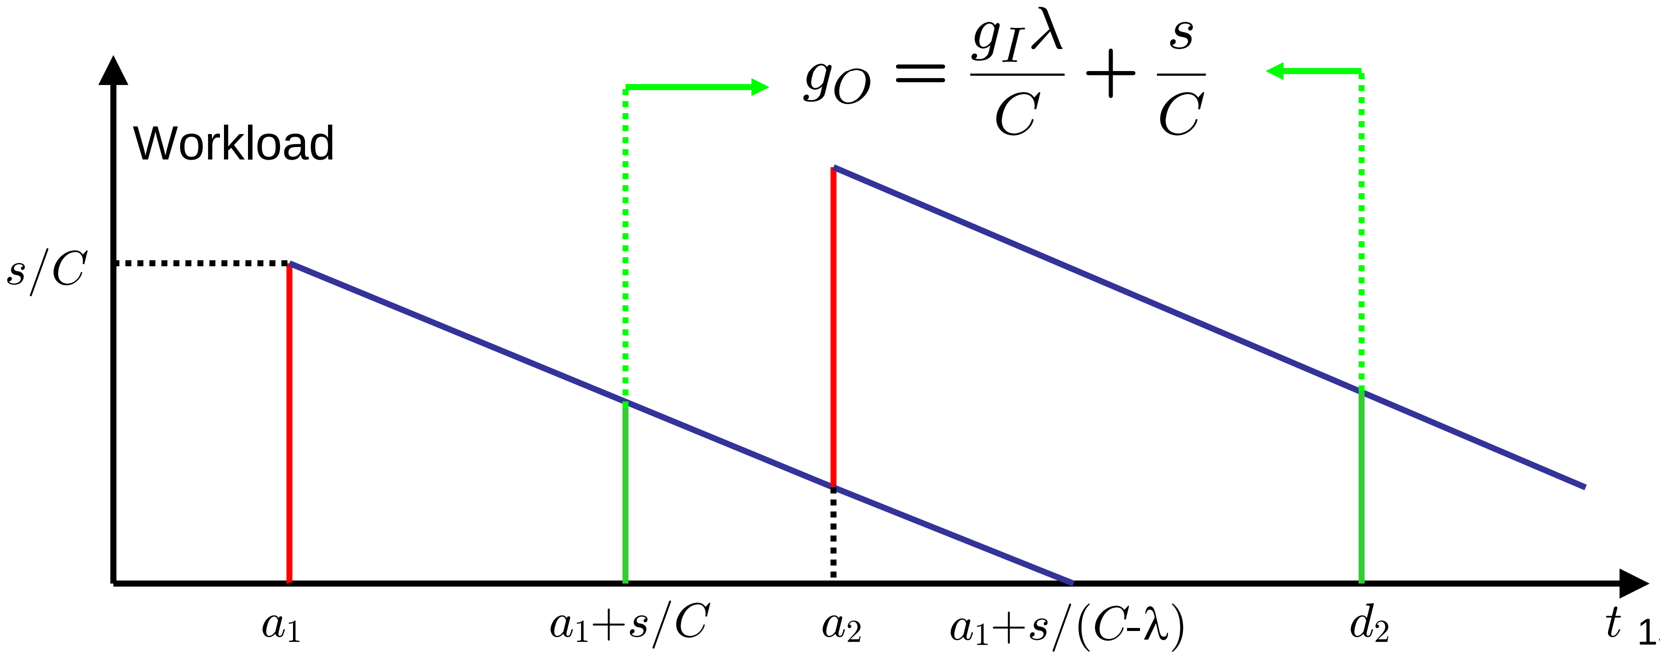

Background: Single-Hop Fluid Curves

- Single-hop fluid rate response curves:

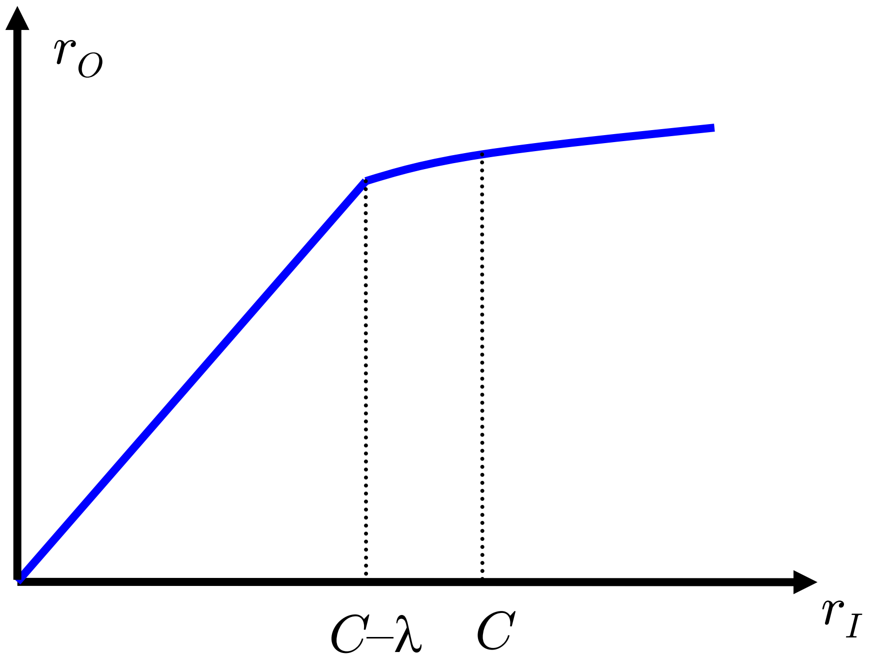

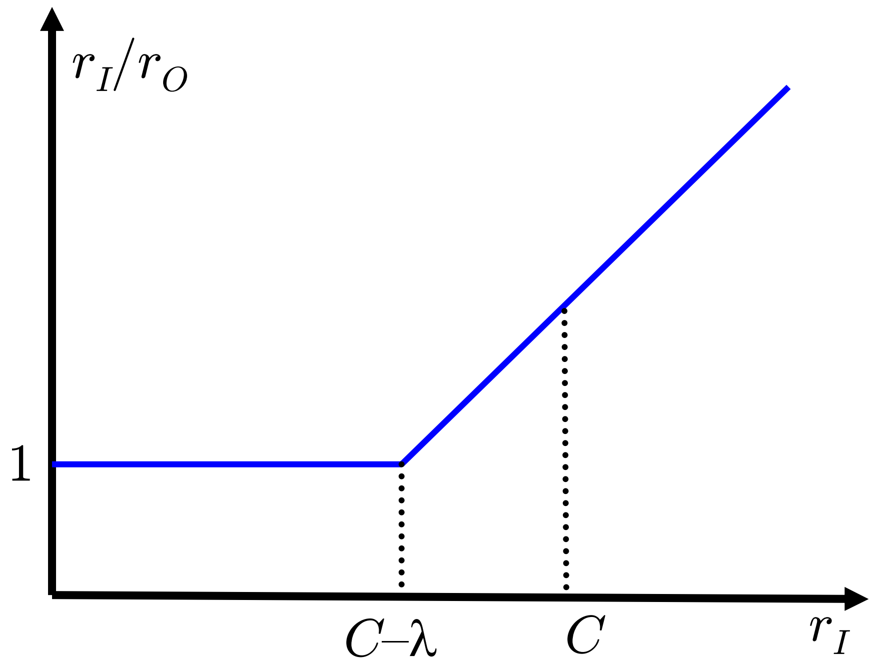

$$r_O = \min\left(r_I, \frac{r_I}{r_I + \lambda} C\right).$$

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

Background: Single-Hop Fluid Curves

- Existing techniques are based on the single-hop fluid curves

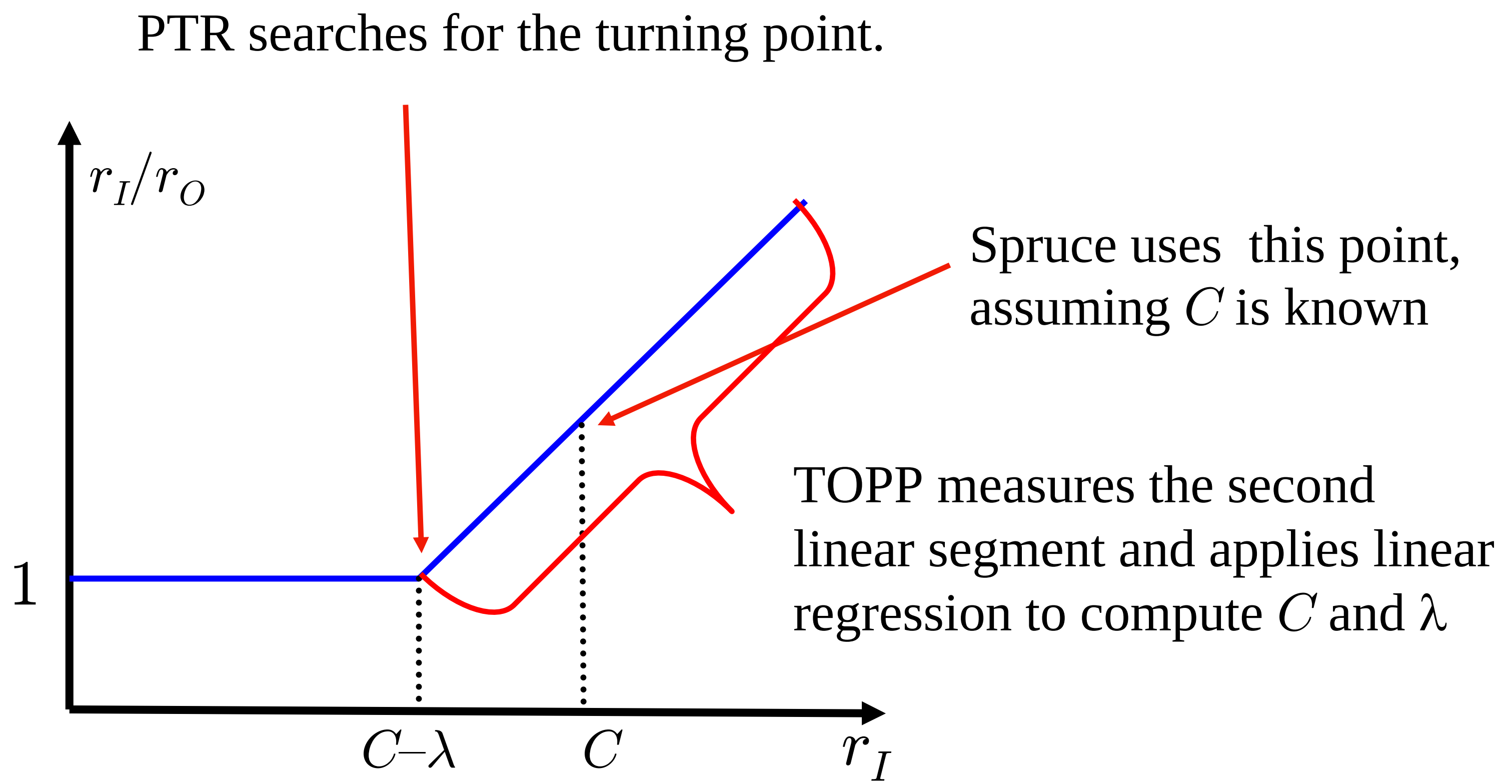

Background: Limitations

- Major limitations of single-hop fluid models
  - Gives no proof that the model also applies to bursty cross-traffic
  - Ignores the impact of non-tight links
  - Provides no insights on the impact of packet train parameters on measurement accuracy

- Existing techniques were observed to produce 100% errors without knowing why

- A stochastic foundation is needed
  - To address these issues
  - To better understand the sources of measurement errors in the current techniques

Outline

- Background
  - What is available bandwidth and why measure it ?
  - Single-hop fluid model and existing techniques
- Stochastic Foundation
  - Single-Hop Case
  - Multi-Hop Case
- Experimental Verification
- Implications and Future Work

Single-Hop Case: Goal

- Derive the single-hop response curve using a packet-level bursty cross-traffic arrival

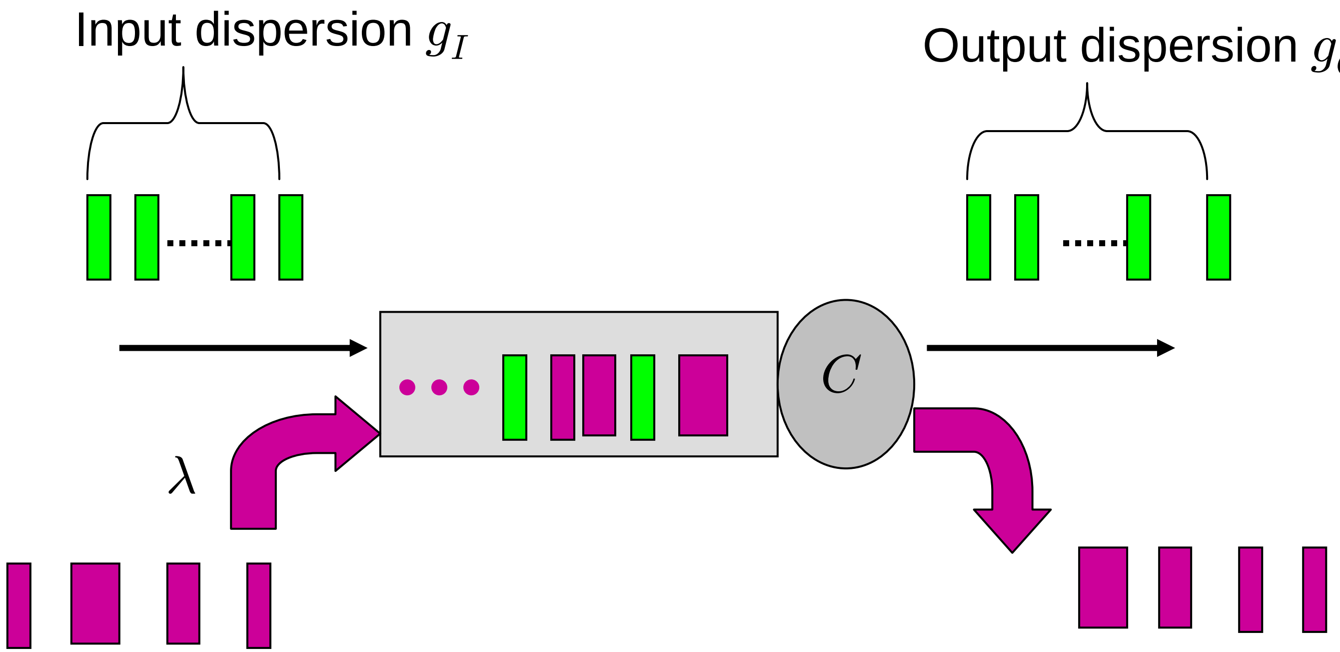

Single-Hop Case: Adapting Fluid Models

- Previous use of fluid models in bursty cross-traffic

$$E[g_O] = \begin{cases} \frac{s+g_I\lambda}{C} & g_I \leq \frac{s}{C-\lambda} \\ g_I & g_I \geq \frac{s}{C-\lambda} \end{cases}$$

- In bursty cross-traffic, g_O varies, take the statistical average of g_O as the output dispersion

- In bursty cross-traffic, traffic arrival rate also varies, interpret λ as the long-term average arrival rate

- However, even with this adaptation, we show that the fluid model is not valid in general

Single-Hop Case: Stochastic Curve

- Stochastic Response Curve
  - We present the following gap response curve in bursty cross-traffic:

$$E[g_O] = \frac{g_I\lambda + s}{C} + \frac{E[I]}{n-1}$$
$$= g_I + \frac{E[R]}{n-1}.$$

  - The two additional terms do not show up in fluid traffic, but do have an effect in bursty cross-traffic

Single-Hop Case: What is $E[I]$

- I is the random variable indicating the hop idle time during the packet-train arriving interval

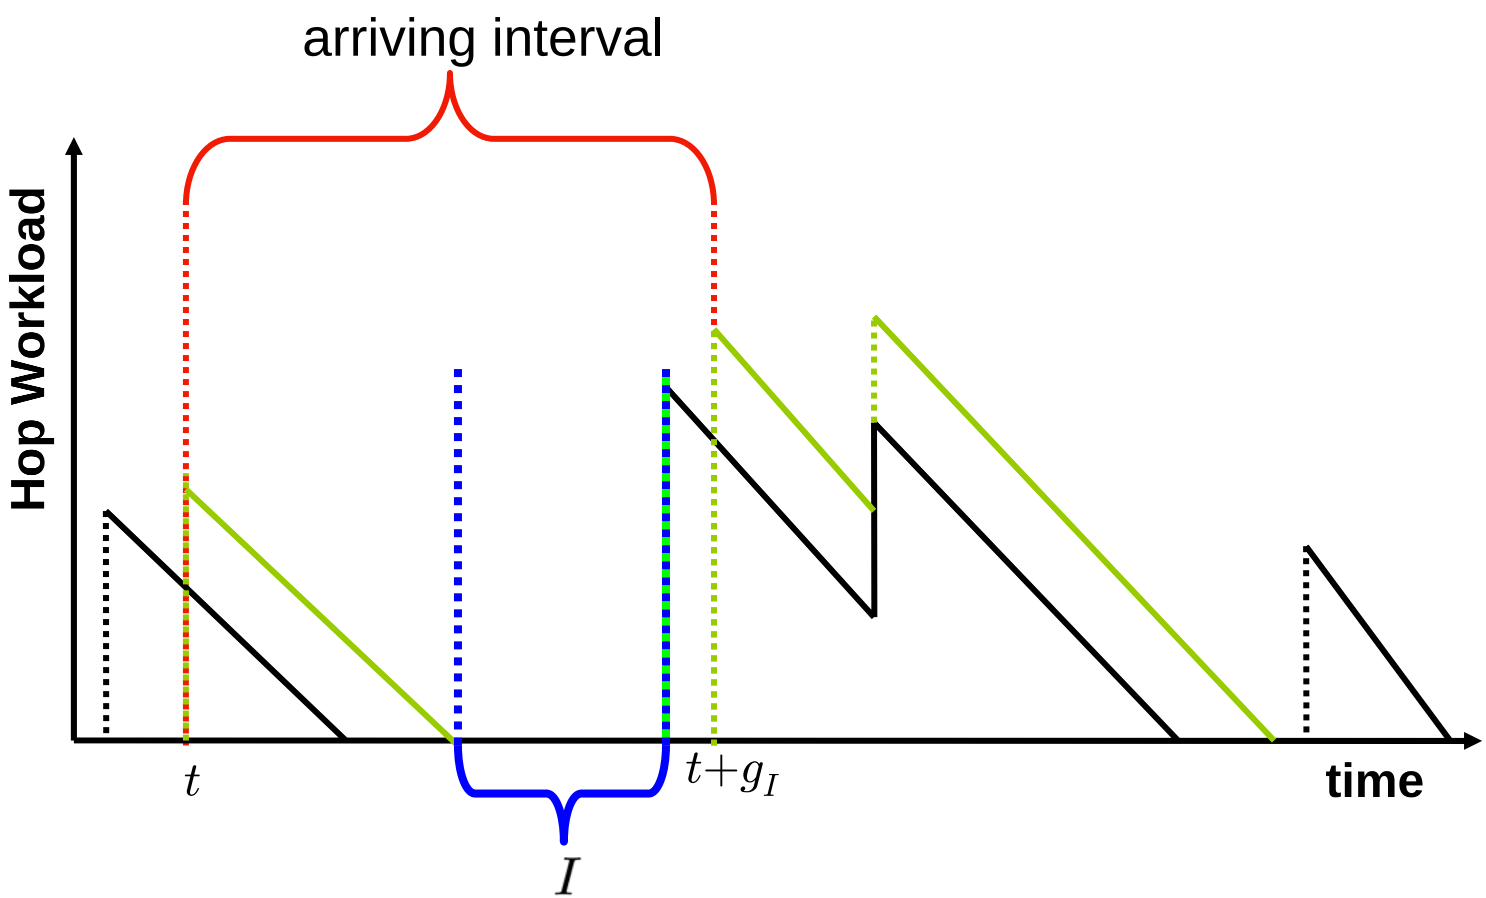

Single-Hop Case: What is $E[R]$

- R is the extra queuing delay imposed on the last packet by the preceding packets in the same probing train

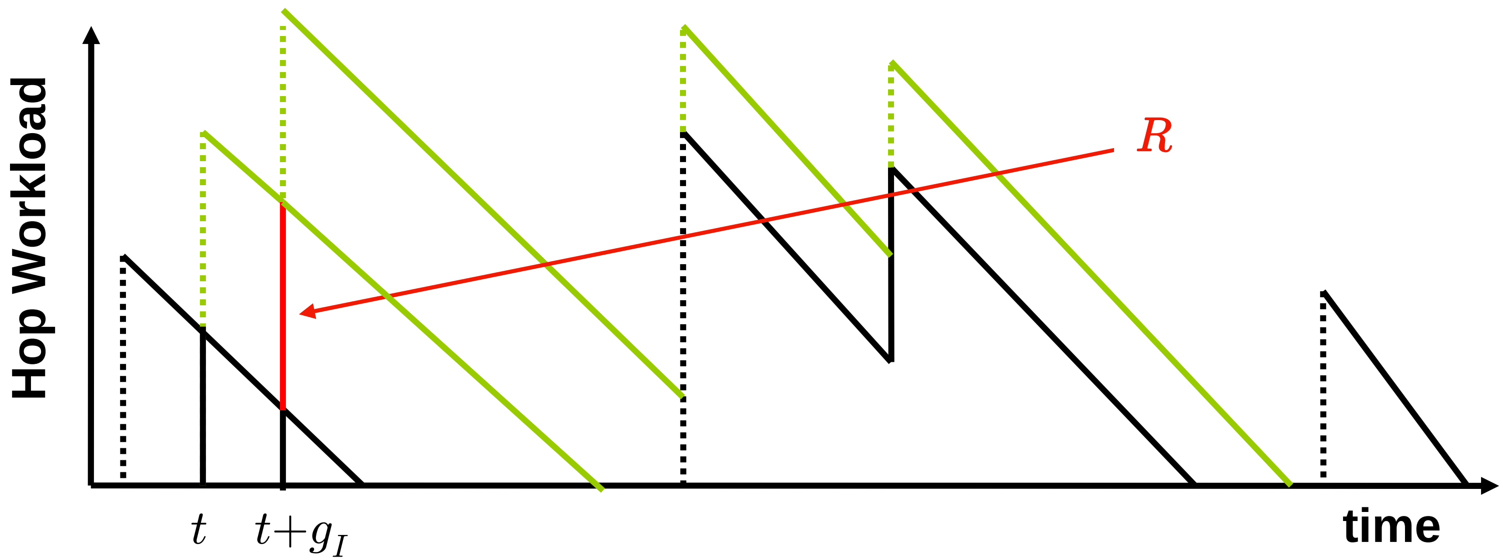

Single-Hop Case: Intuitive Interpretation

- The two expressions describe $E[g_O]$ from two different angles

$$E[g_O] = \frac{g_I \lambda + s}{C} + \frac{E[I]}{n-1}$$

$$= g_I + \frac{E[R]}{n-1}$$

Hop activities between the departures of the pair

What causes the difference between the input and output dispersion?

Single-Hop Case: Response Deviation 1

- The two additional terms cause the stochastic response curve to deviate from the fluid curve

$$E[g_O] - \max\left(g_I, \frac{g_I\lambda + s}{C}\right) = \begin{cases} \dfrac{E[I]}{n-1} & r_I \geq C - \lambda \\ \dfrac{E[R]}{n-1} & r_I \leq C - \lambda \end{cases}$$

Response Deviation β

Single-Hop Case: Response Deviation 2

- Response Deviation as a Function of g_I

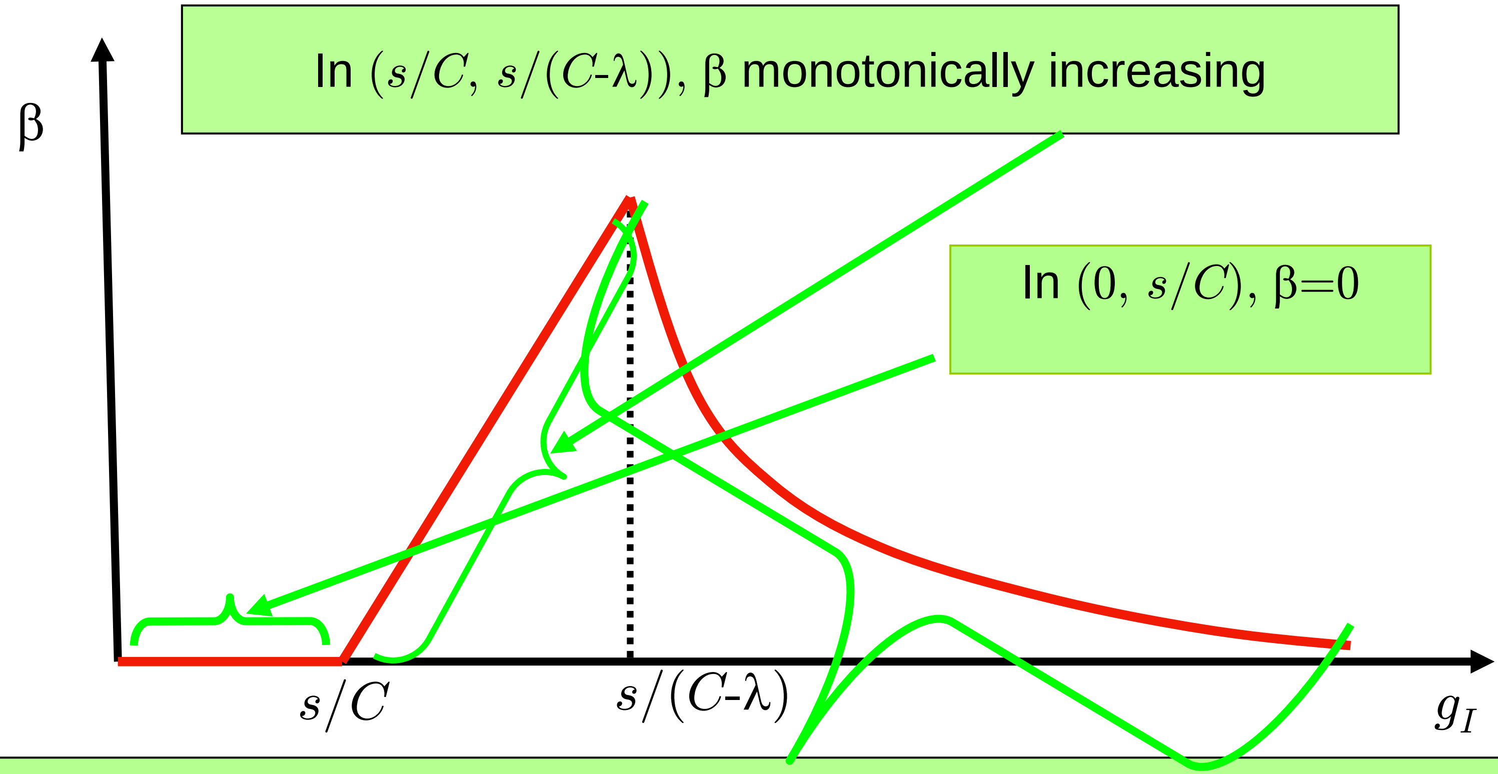

When $g_I > s/(C\text{-}\lambda)$, β monotonically decreases and asymptotically converges to 0.

Single-Hop Case: Response Deviation 3

- Transformed rate response curve

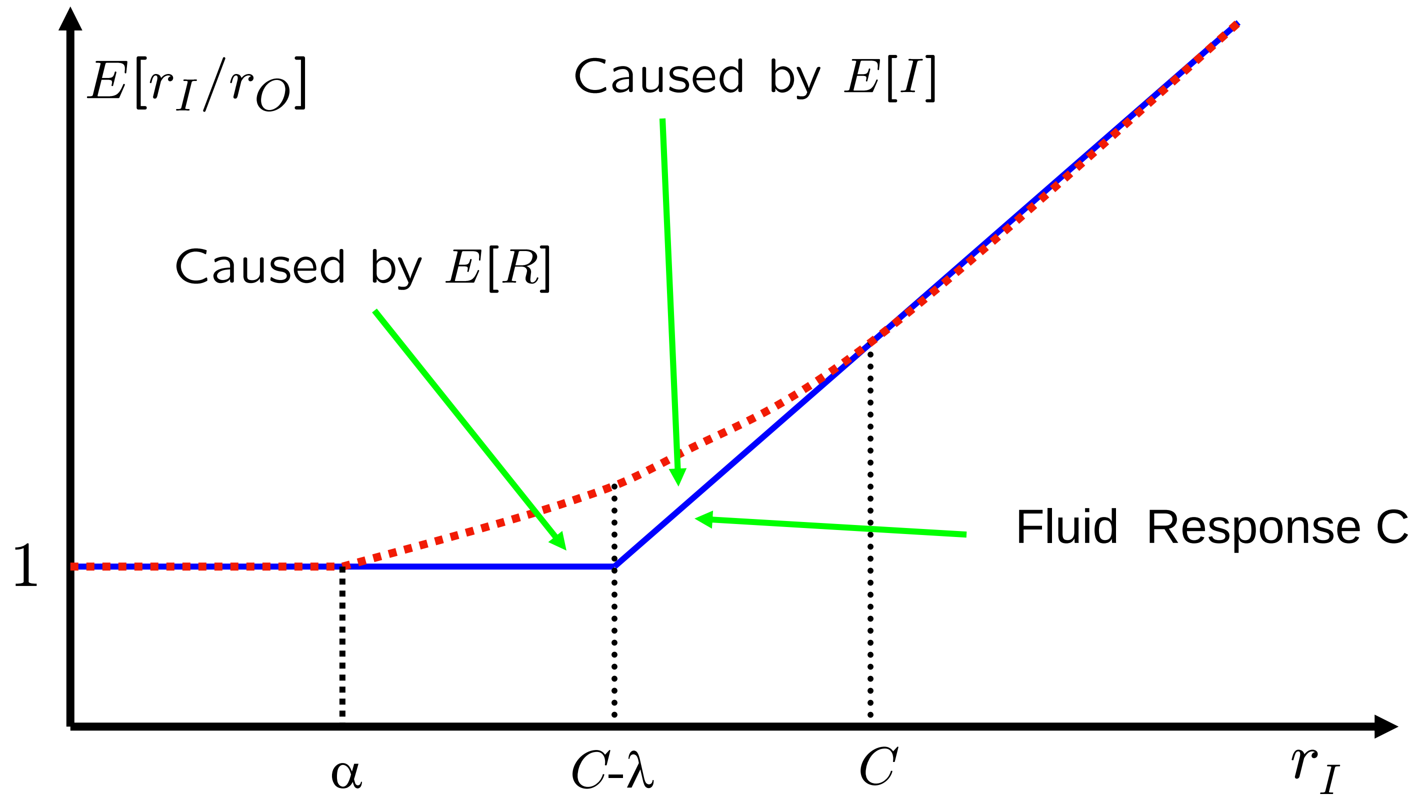

Single-Hop Case: Packet-Train Parameters

- When packet train length n increases, response deviation vanishes
- Consider a packet-train of infinite length
  - When $r_I < C - \lambda$, queue is stable, mean queuing delay is bounded, so is the extra queuing delay term

$$E[g_O] = g_I + \frac{E[R]}{n-1}$$

  - When $r_I > C - \lambda$, queue goes unbounded, the amount of hop idle time is bounded

$$E[g_O] = \frac{g_I \lambda + s}{C} + \frac{E[I]}{n-1}$$

Single-Hop Case: Summary

- Two additional terms show up in bursty cross-traffic, causing the single-hop stochastic curve to deviate from the fluid curve
- As packet-train length n increases, the stochastic curve approaches the fluid curve
- Conclusion: The fluid curve is a valid first-order approximation of the stochastic curve only when packet-train length is sufficiently large

Outline

- Background
  - what is available bandwidth and why measure it ?
  - deterministic fluid foundation and existing techniques
- Stochastic Foundation
  - Single-Hop Case
  - Multi-Hop Case
- Experimental Verification
- Implications and Future Work

Multi-Hop Case: Problem Statement

- An N-hop path probed by packet trains of length n

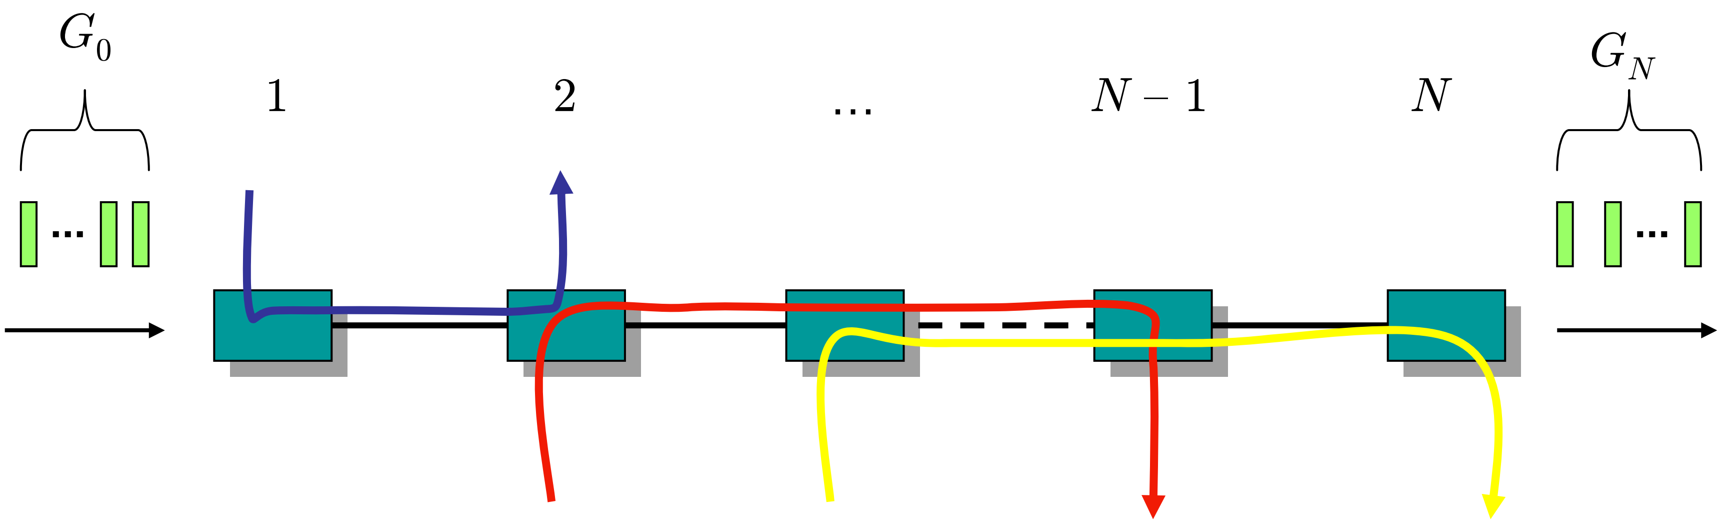

- **Goal**: understand the relationship between $E[G_N]$ and G_0 under arbitrary cross-traffic
  – That is the multi-hop probing response curve

Multi-Hop Case: Simplest Settings

- Consider fluid cross-traffic with one-hop persistent routing

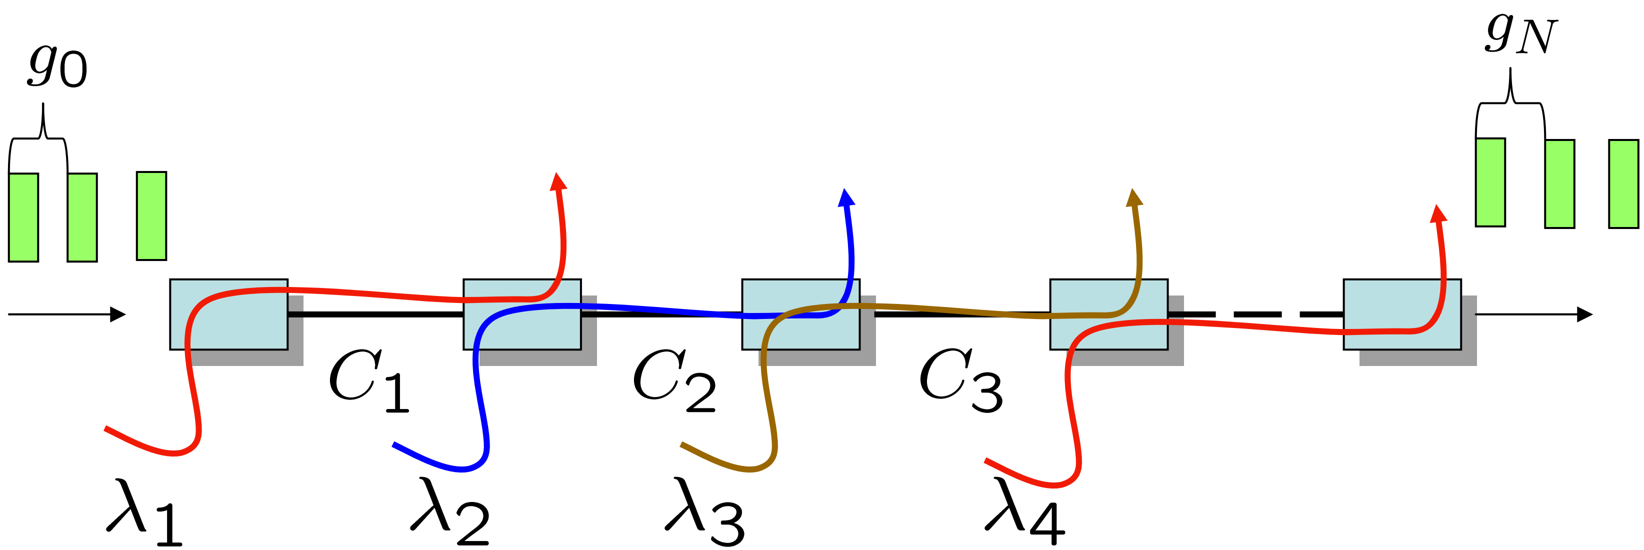

- Mathematically:

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

Multi-Hop Case: Relaxing Fluid Constraint

- When relaxing the fluid constraint, we get

$$E[G_i] = \frac{E[G_{i-1}]\lambda_i + s}{C_{i-1}} + \frac{E[I_i]}{n-1}$$
$$= E[G_{i-1}] + \frac{E[R_i]}{n-1}.$$

- The response deviation at link i is

$$\beta_i = E[G_i] - g_i = \begin{cases} \beta_{i-1} + \frac{E[R_i]}{n-1} & g_i = g_{i-1} \\ \frac{\beta_{i-1}\lambda_i}{C_i} + \frac{E[I_i]}{n-1} & g_i > g_{i-1} \end{cases}$$

Multi-Hop Case: Effect of Packet-Train Length

- In one-hop persistent routing, It is easy to show using induction that

$$\beta_i \longrightarrow 0, \quad \text{as} \quad n \longrightarrow \infty$$

- Conclusion: Multi-Hop curve approaches from above its fluid counterpart as packet-train length increases
  - This result also applies to arbitrary CT routing
  - Similar reasons, complex math description

Multi-Hop Case: Fluid Response Curves

- Multi-Hop Fluid Curves for Different CT Routing

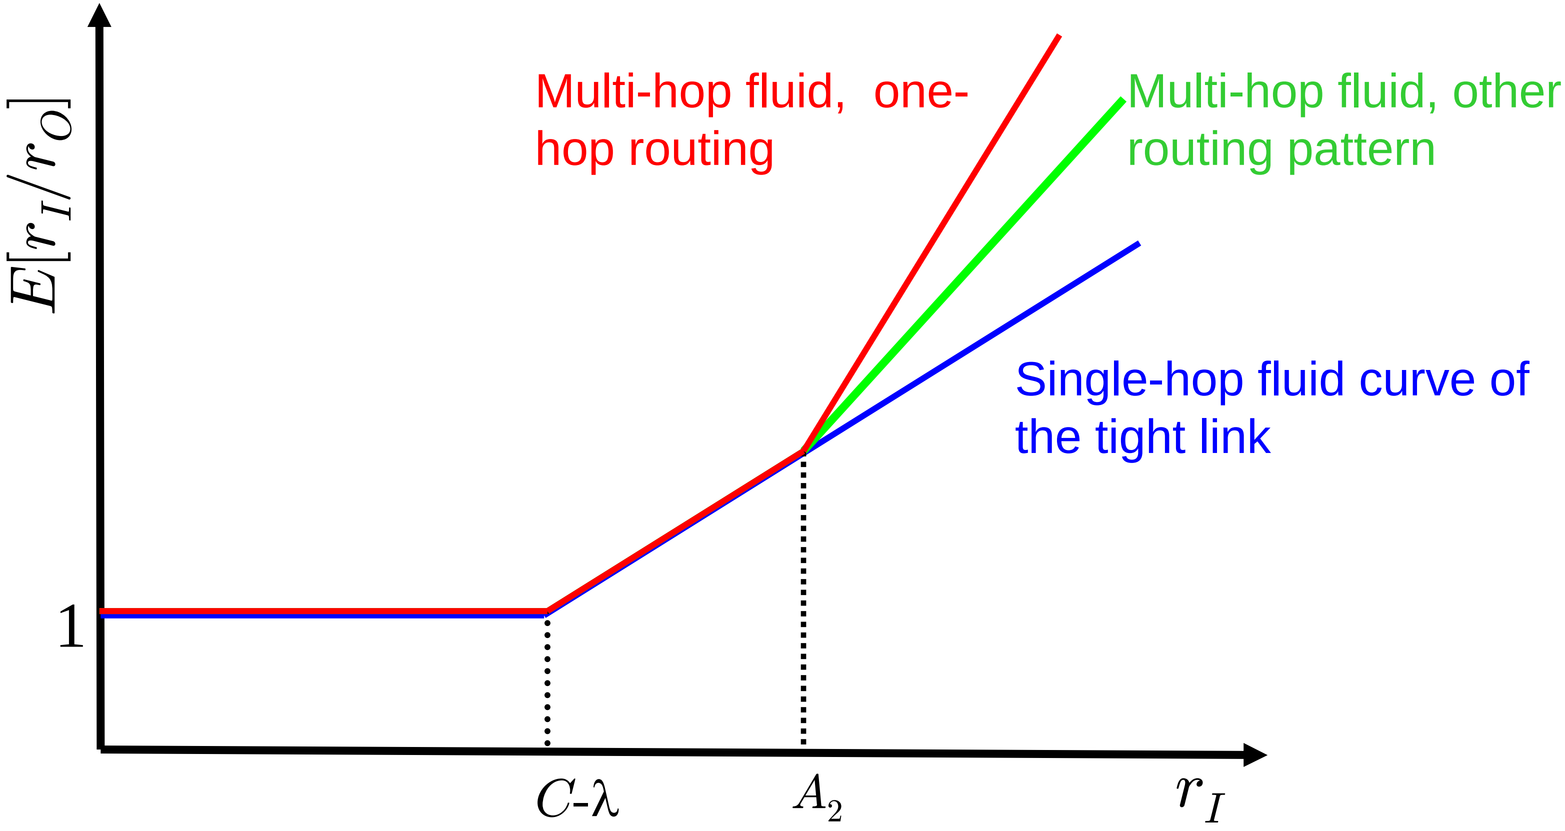

Multi-Hop Case: Stochastic Response Curves

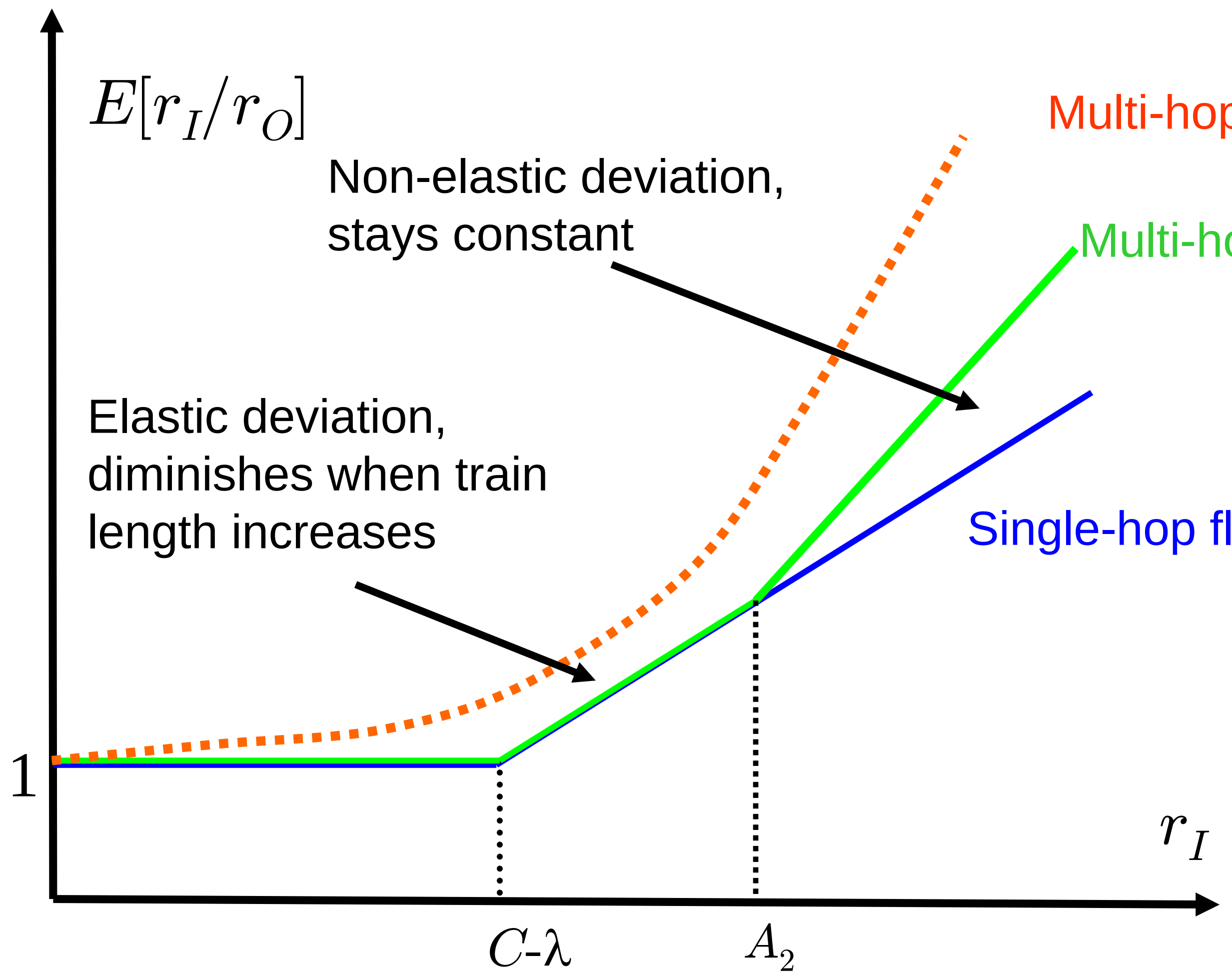

Multi-Hop Case: Stochastic Response Curves

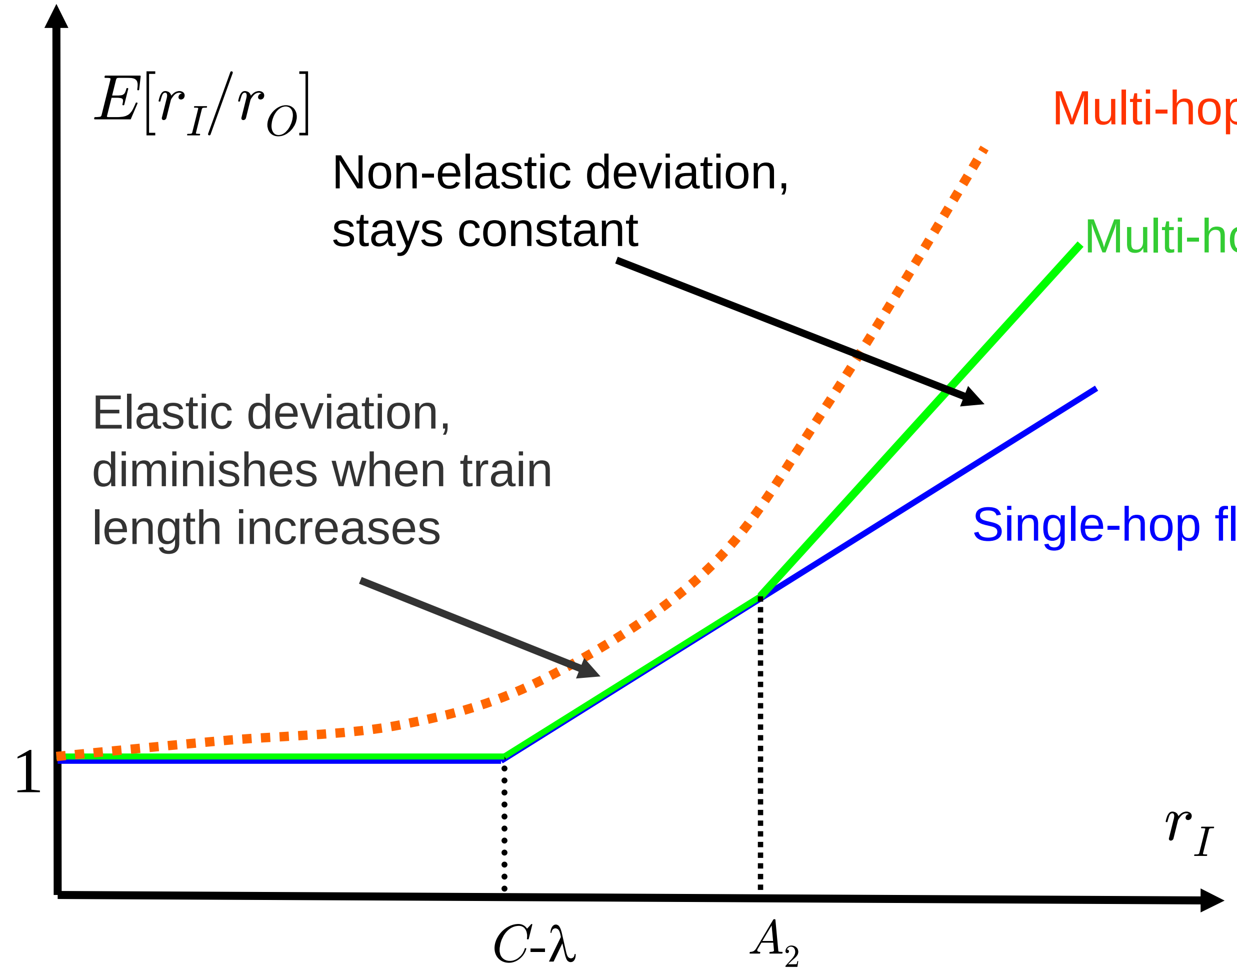

Multi-Hop Case: Stochastic Response Curves

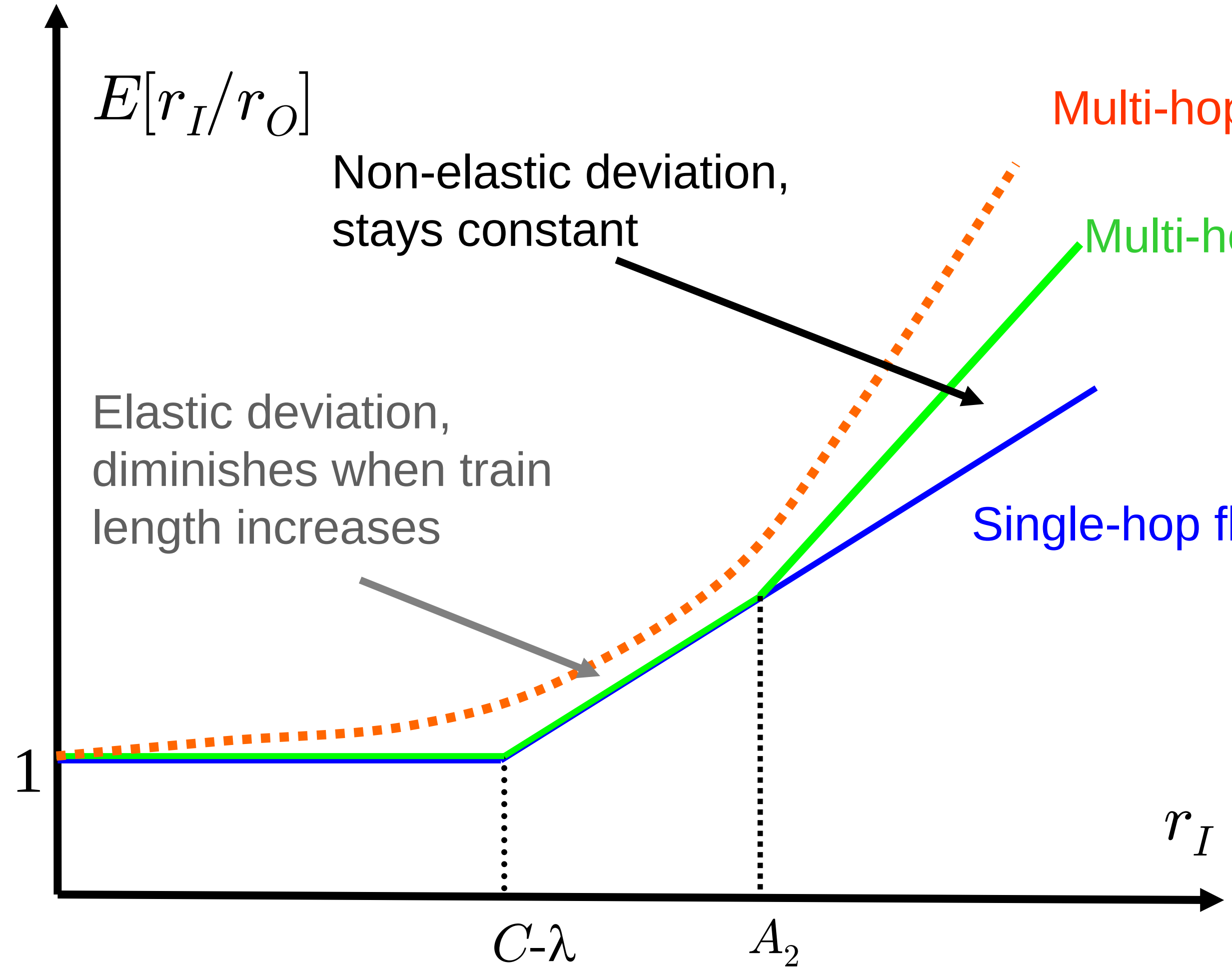

Multi-Hop Case: Stochastic Response Curves

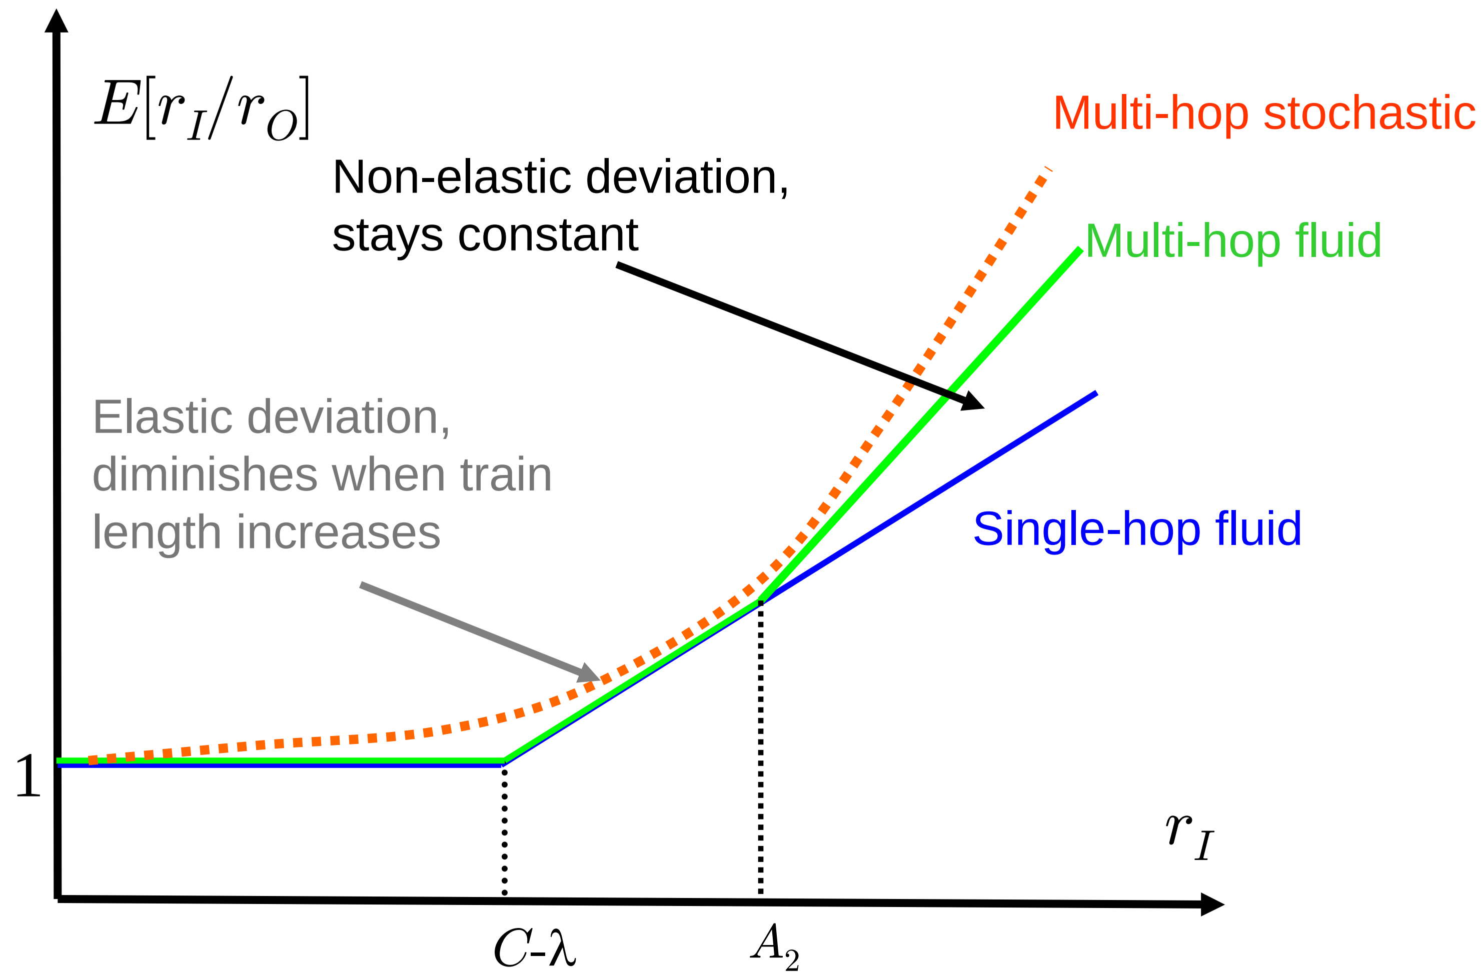

Multi-Hop Case: Stochastic Response Curves

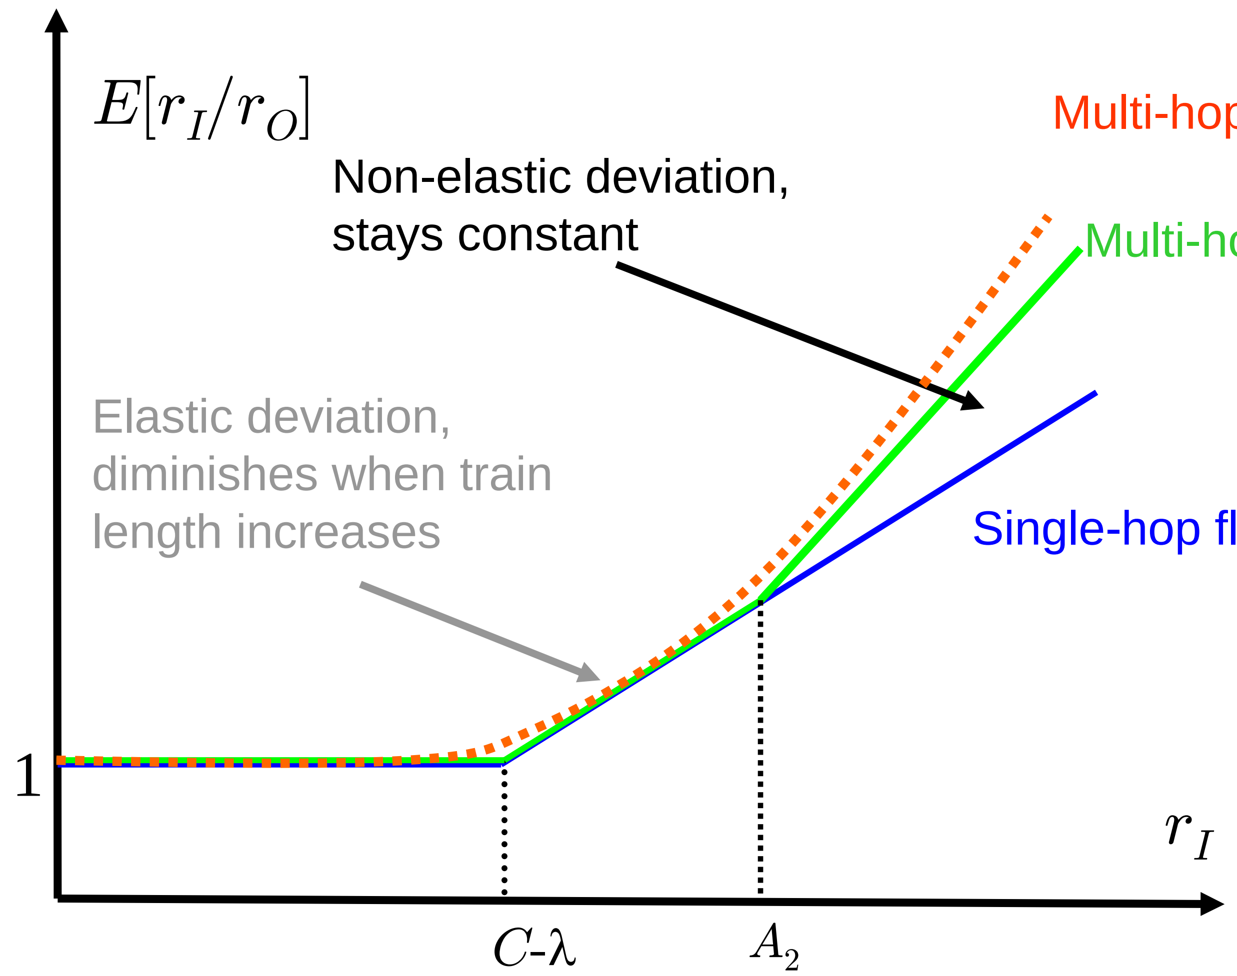

Multi-Hop Case: Stochastic Response Curves

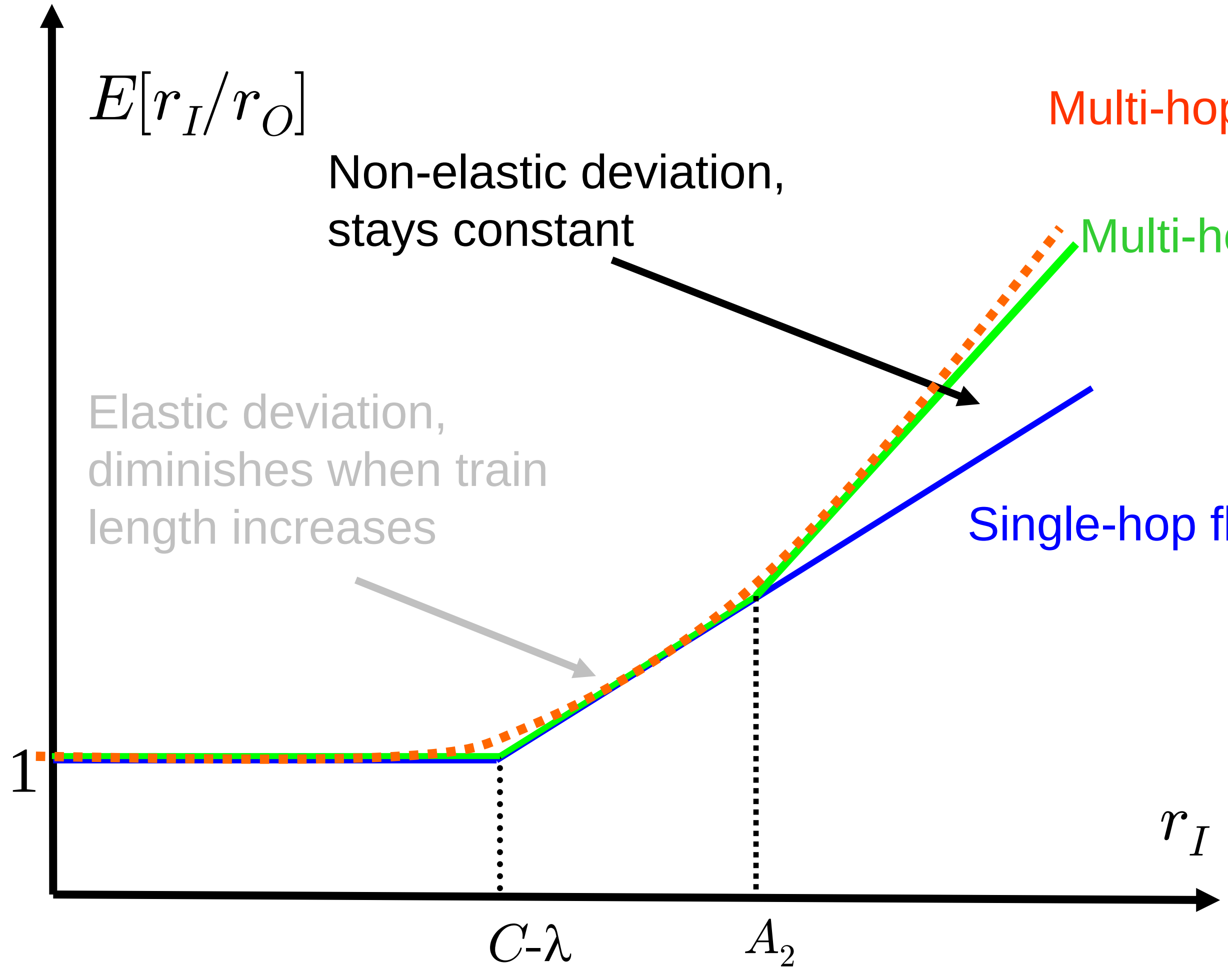

Multi-Hop Case: Stochastic Response Curves

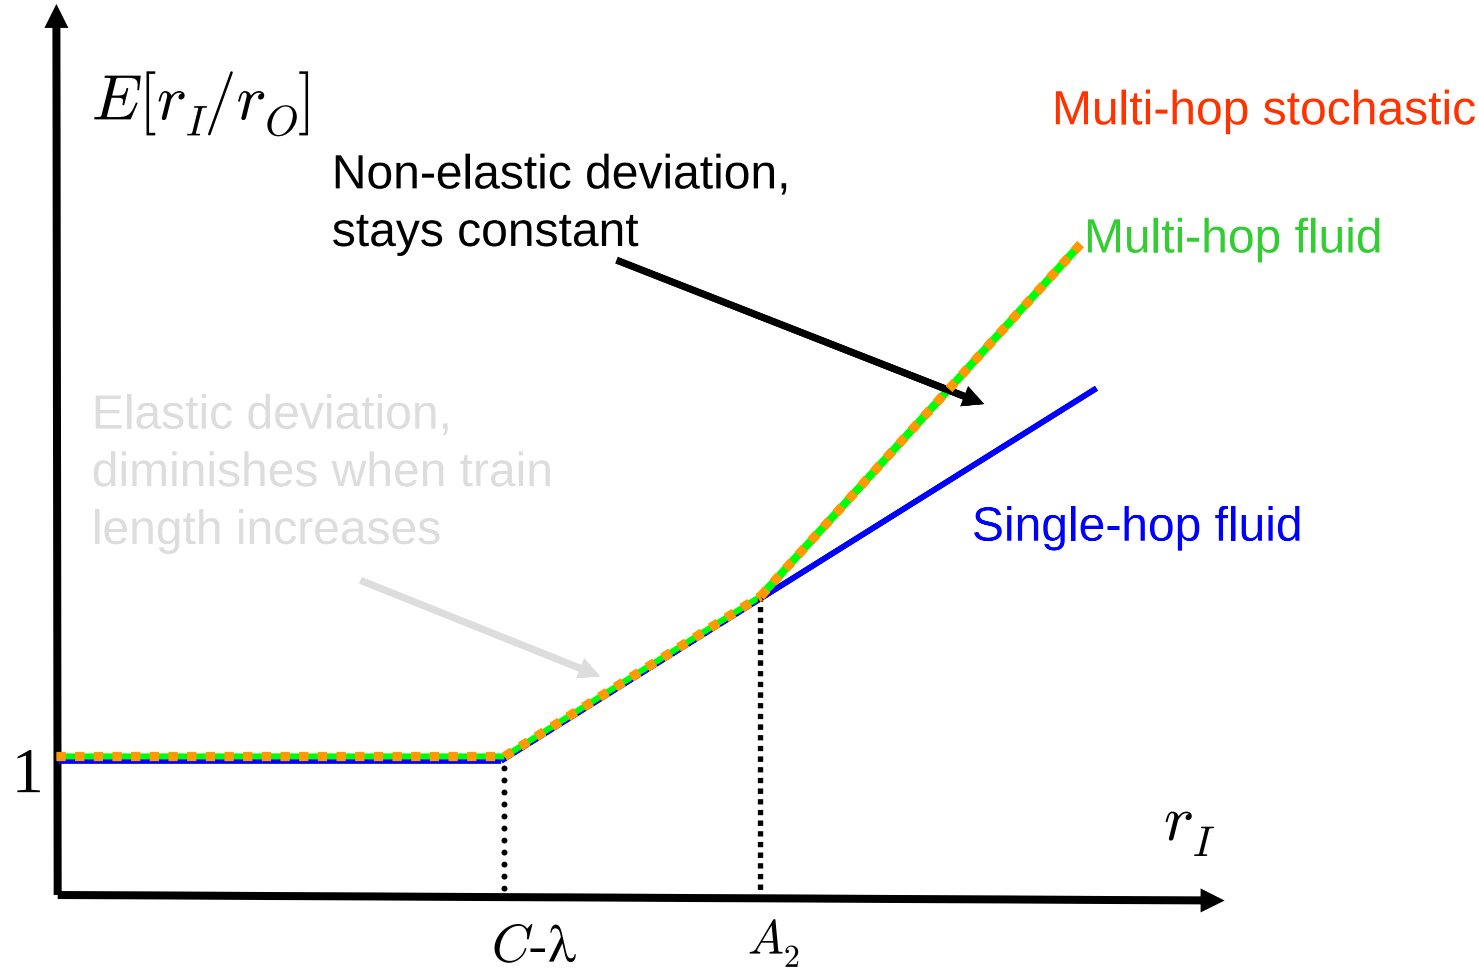

Outline

- Background
  - what is available bandwidth and why measure it ?
  - deterministic fluid foundation and existing techniques
- Stochastic Foundation
  - Single-Hop Case
  - Multi-Hop Case
- Experimental Verification
- Implications and Future Work

Experimental Verification: Roadmap

- Single-Hop Response Curves
  - NS2 Simulation in Poisson traffic
- Multi-Hop Response Curves
  - Emulab testbed experiment with real traffic traces
  - Real Internet measurement over the RON testbed

Experimental Verification: Single-Hop

- Single-hop path 10mb/s capacity with 3mb/s Poisson cross-traffic

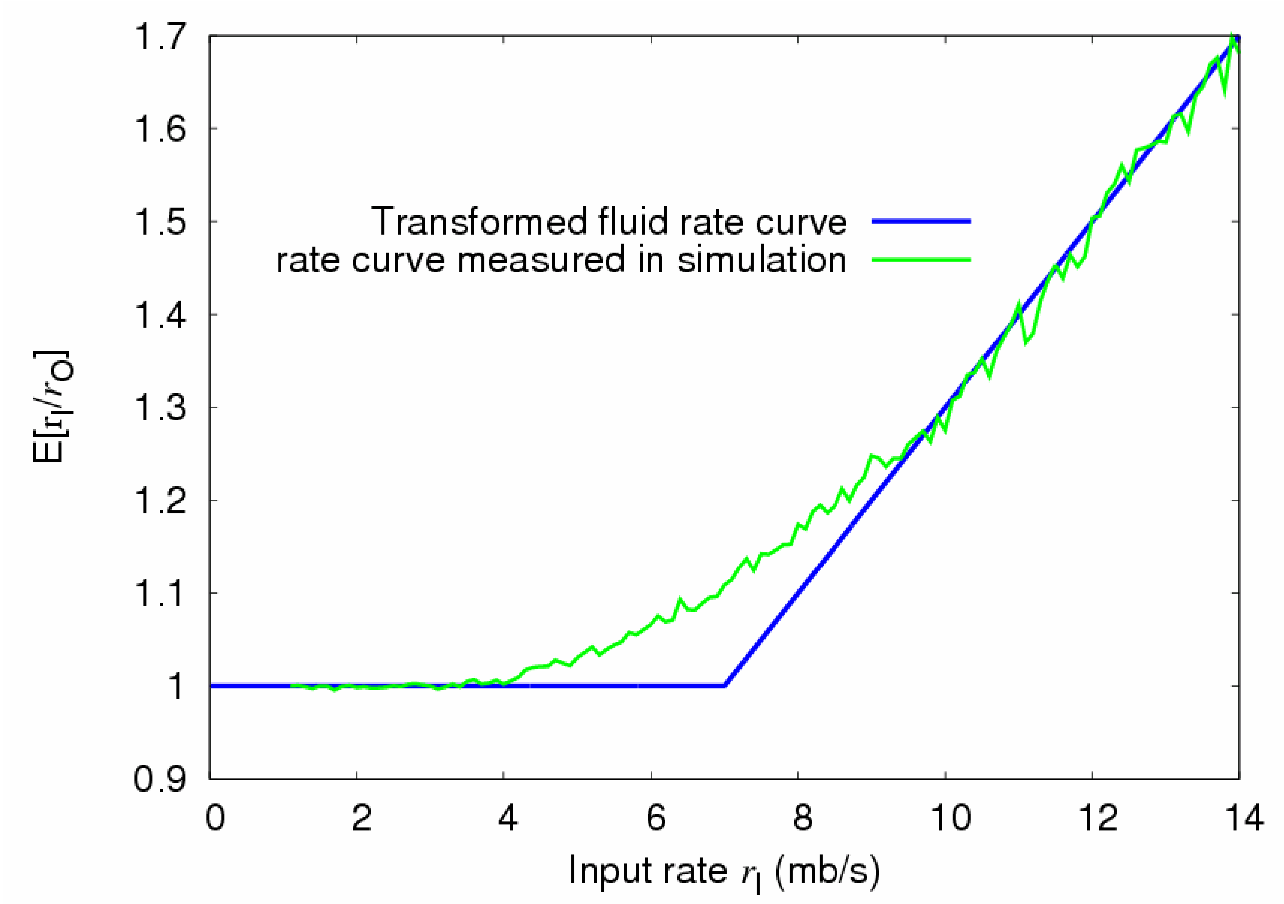

Experimental Verification: Emulab 1

- Emulab Testbed Settings

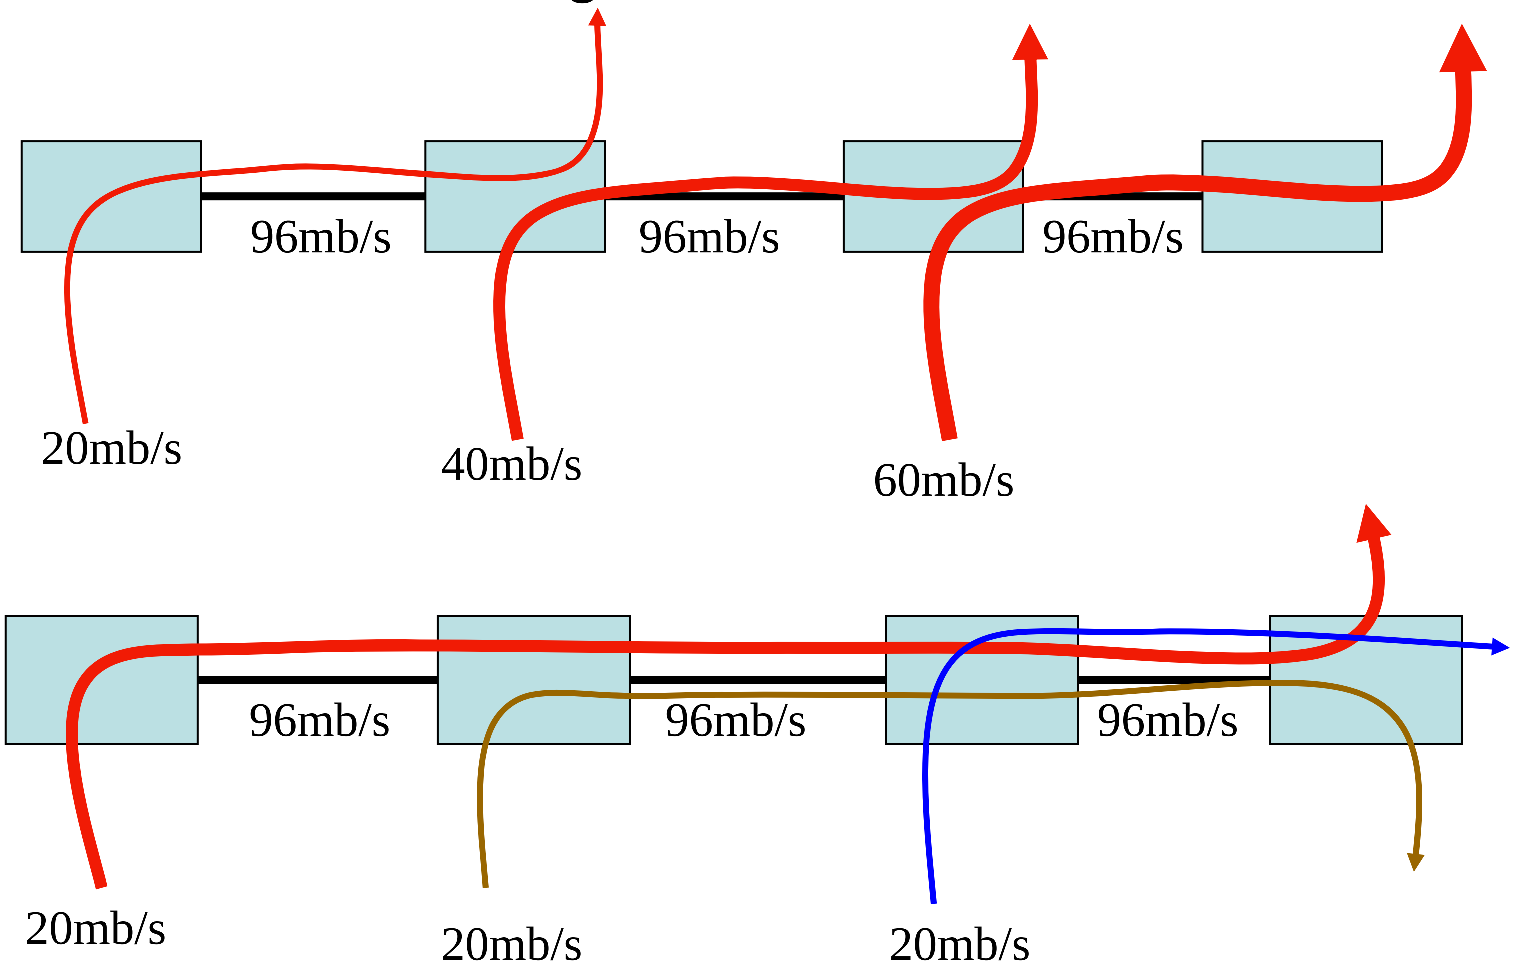

Experimental Verification: Emulab 2

- One-Hop Persistent Routing Case

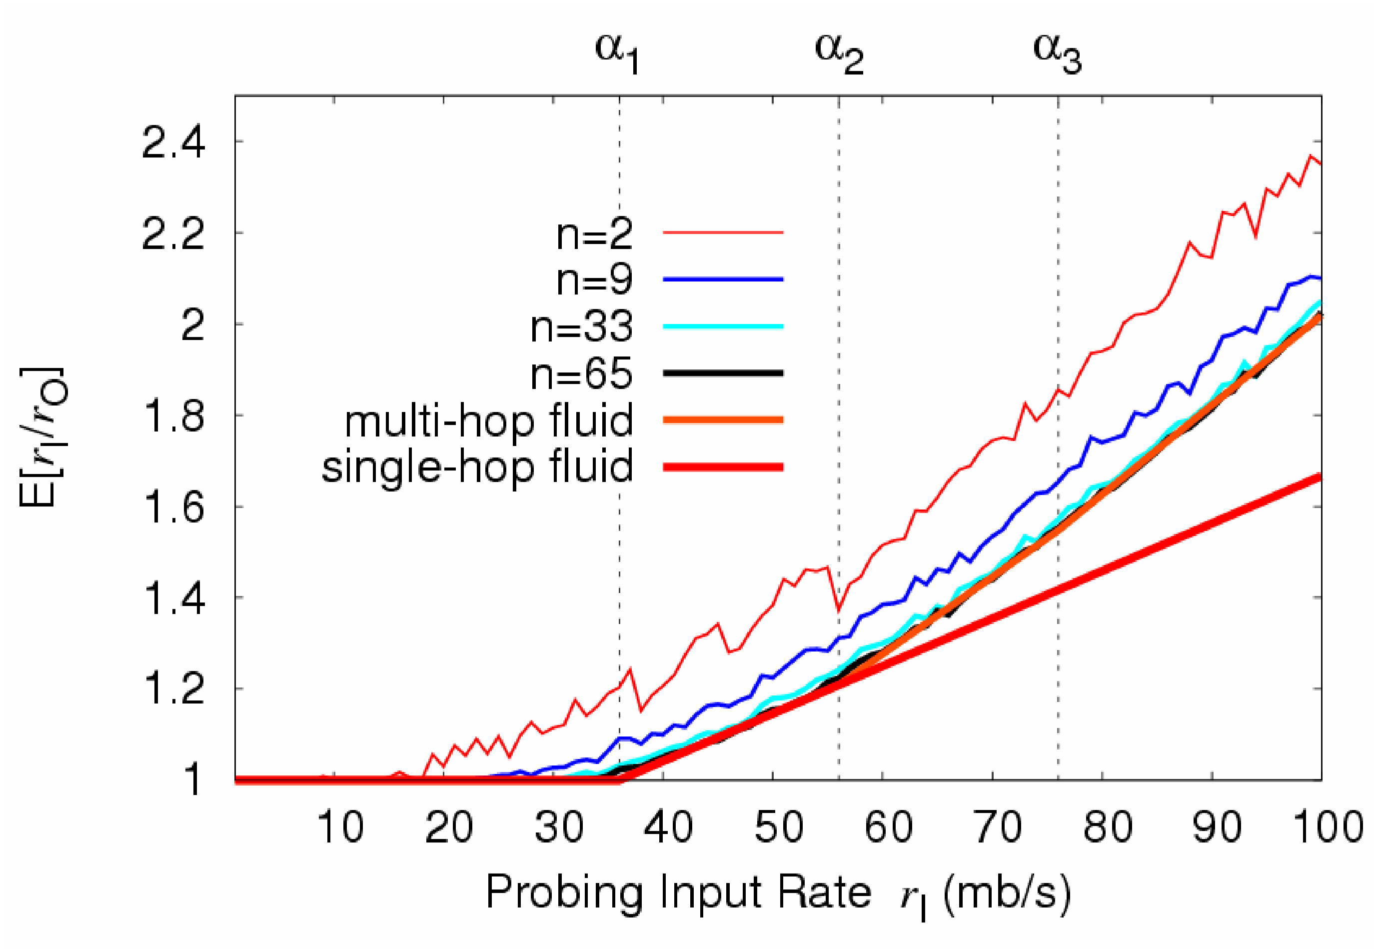

Experimental Verification: Emulab 3

- Path Persistent Routing Case

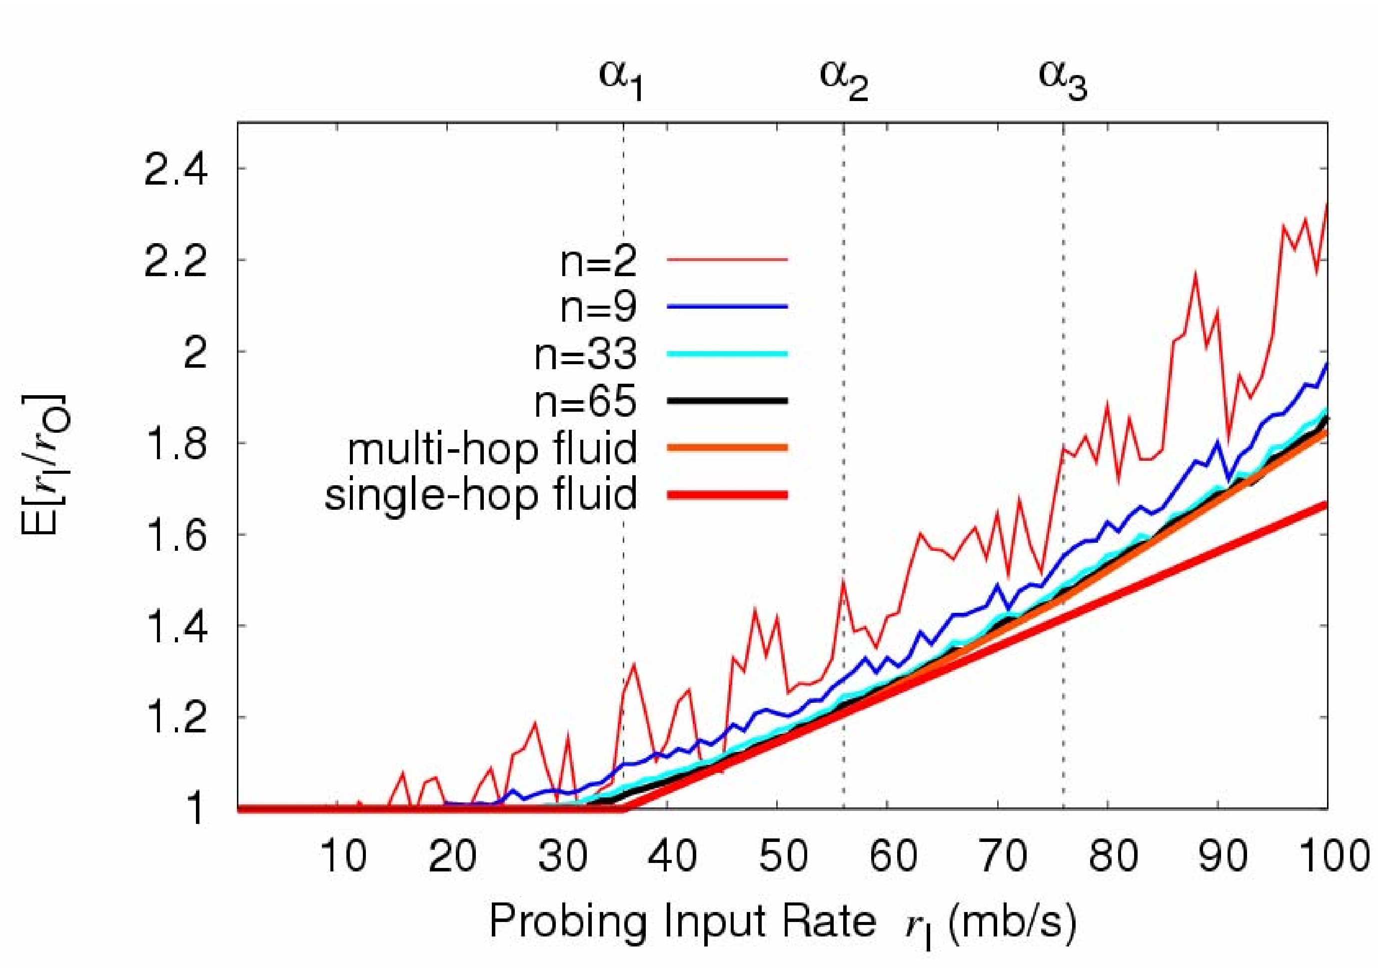

Experimental Verification: Real Internet Data

- We measure the rate response curves for 272 Internet paths over the RON testbed
- Parameters:
  - Input rates: from 10 to 150 mb/s with step 5 mb/s
  - Packet-train length: 129 packets
  - Packet-size: 1500 bytes
  - For each rate, we use 200 trains to estimate $E[G_N]$
- Experiment durations are 20-100 minutes

Experimental Verification: Internet Data

- Cornell→CMU, 5/25/2005

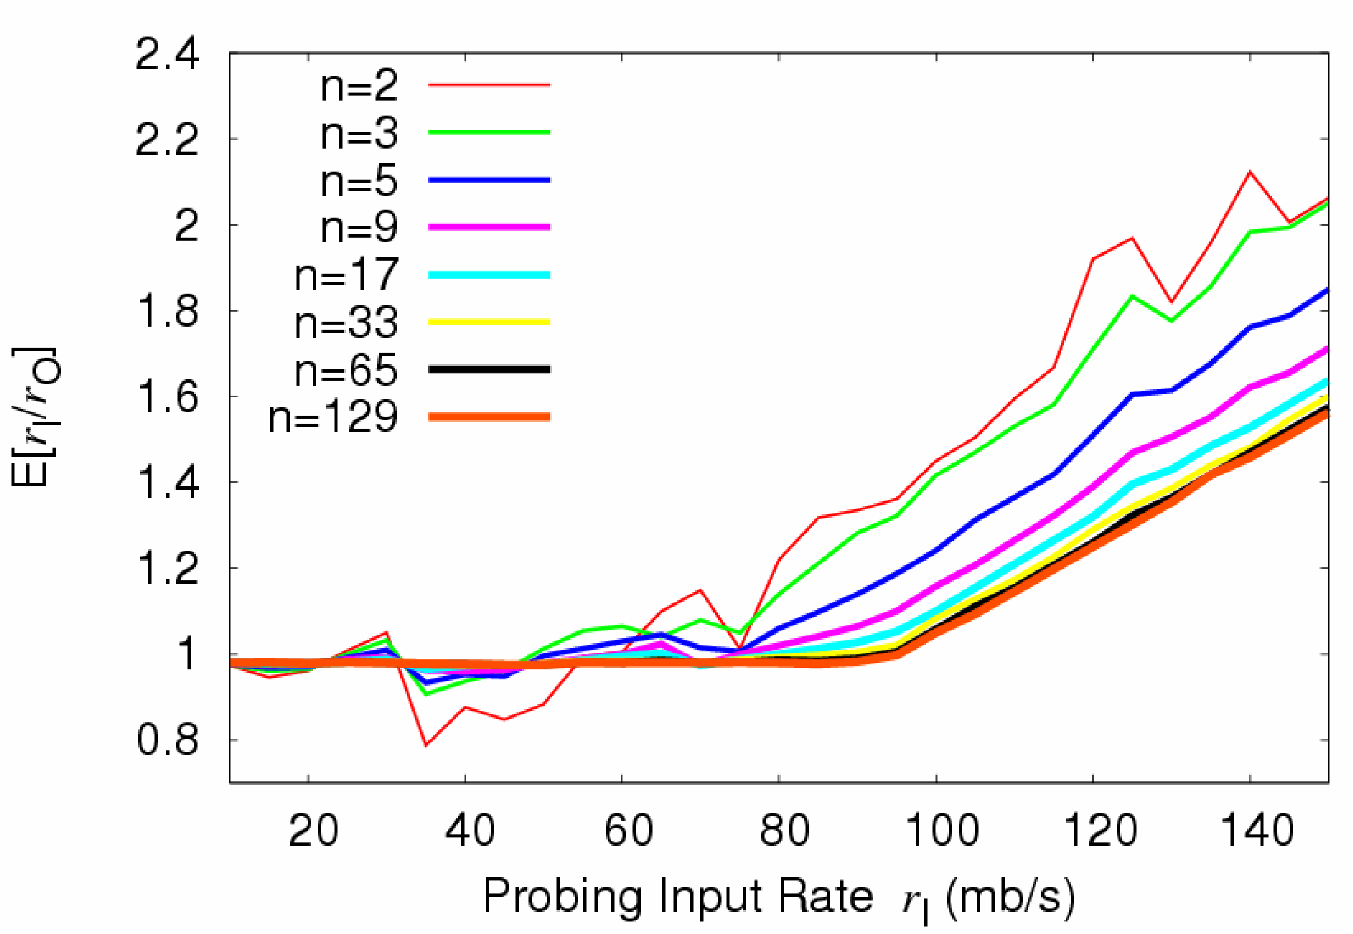

Experimental Verification: Internet Data

- Ana1-gblx → Cornell, 4/29/2005

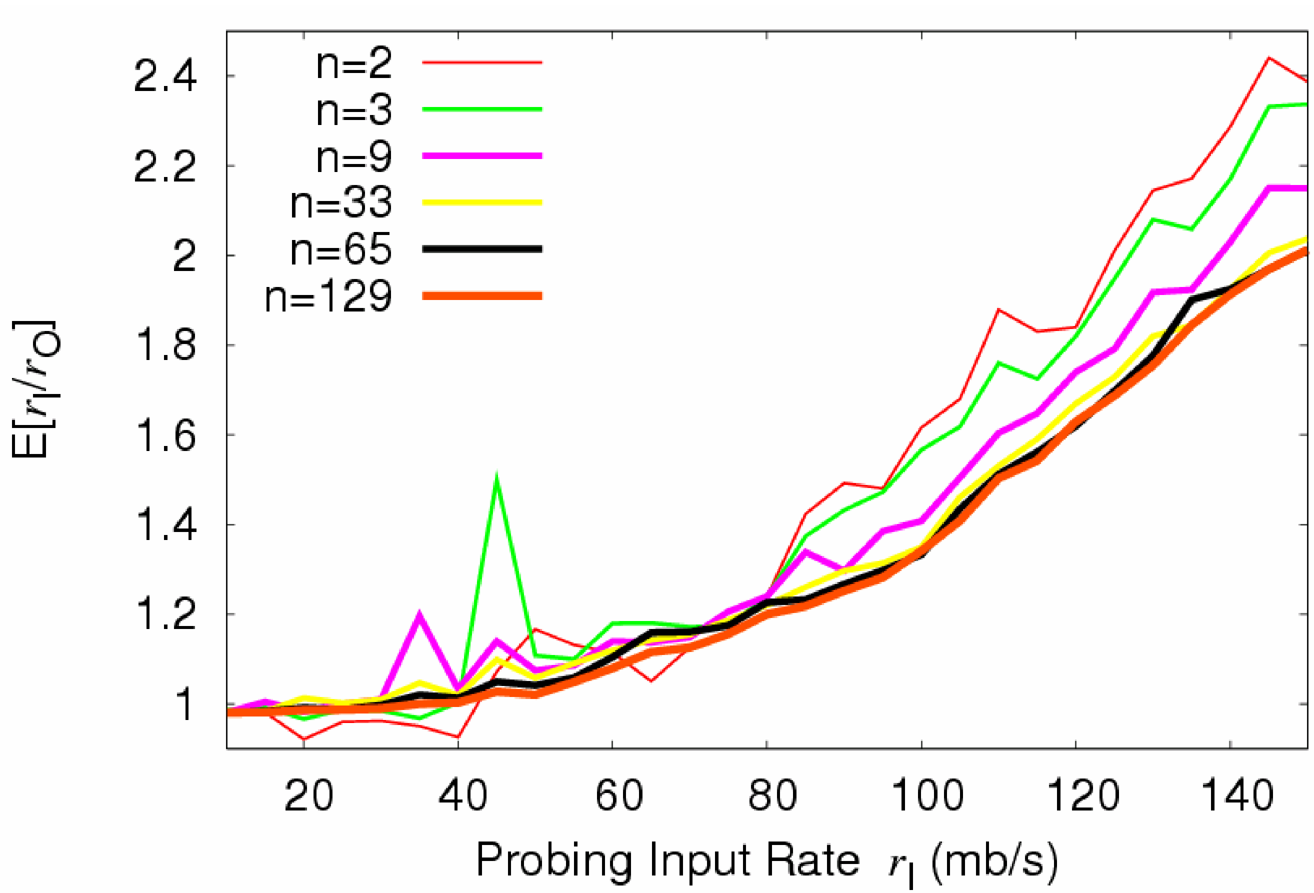

Outline

- Background
  - what is available bandwidth and why measure it ?
  - deterministic fluid foundation and existing techniques
- Stochastic Foundation
  - Single-Hop Case
  - Multi-Hop Case
- Experimental Verification
- Implications

Implications: TOPP

- TOPP uses packet-pairs to measure the stochastic response curve and implicitly assumes that it is the same as the fluid curve
  - Our results show that the two are not the same even for a single-hop path (response deviation)
  - Increasing the packet-train length can reduce the response deviation to a negligible level, and make TOPP work in practice
  - Our Internet measurement shows a length of several tens (e.g., 30) is usually enough

Implications: TOPP

- Use the Poisson simulation as an example

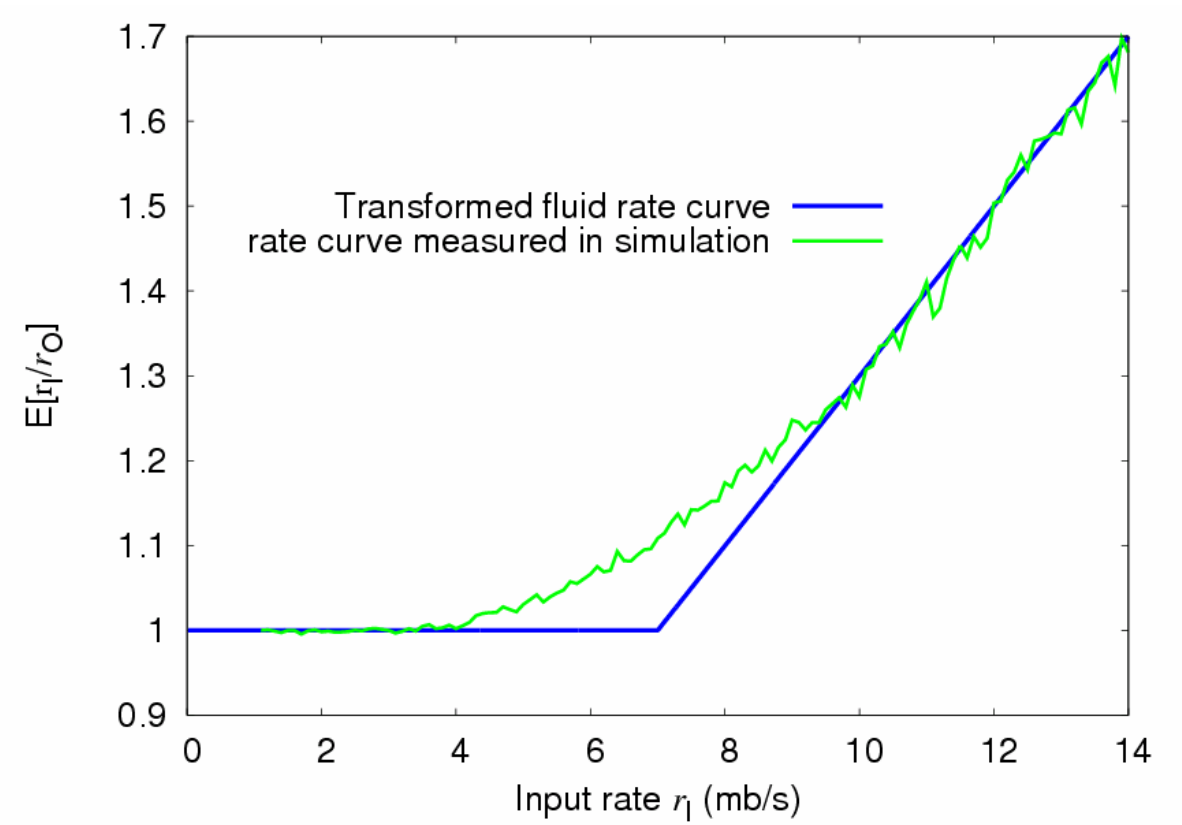

C=10	λ=3	A=7
35.81	32.38	3.43

Implications: Spruce 1

- Spruce is unbiased in a single-hop path

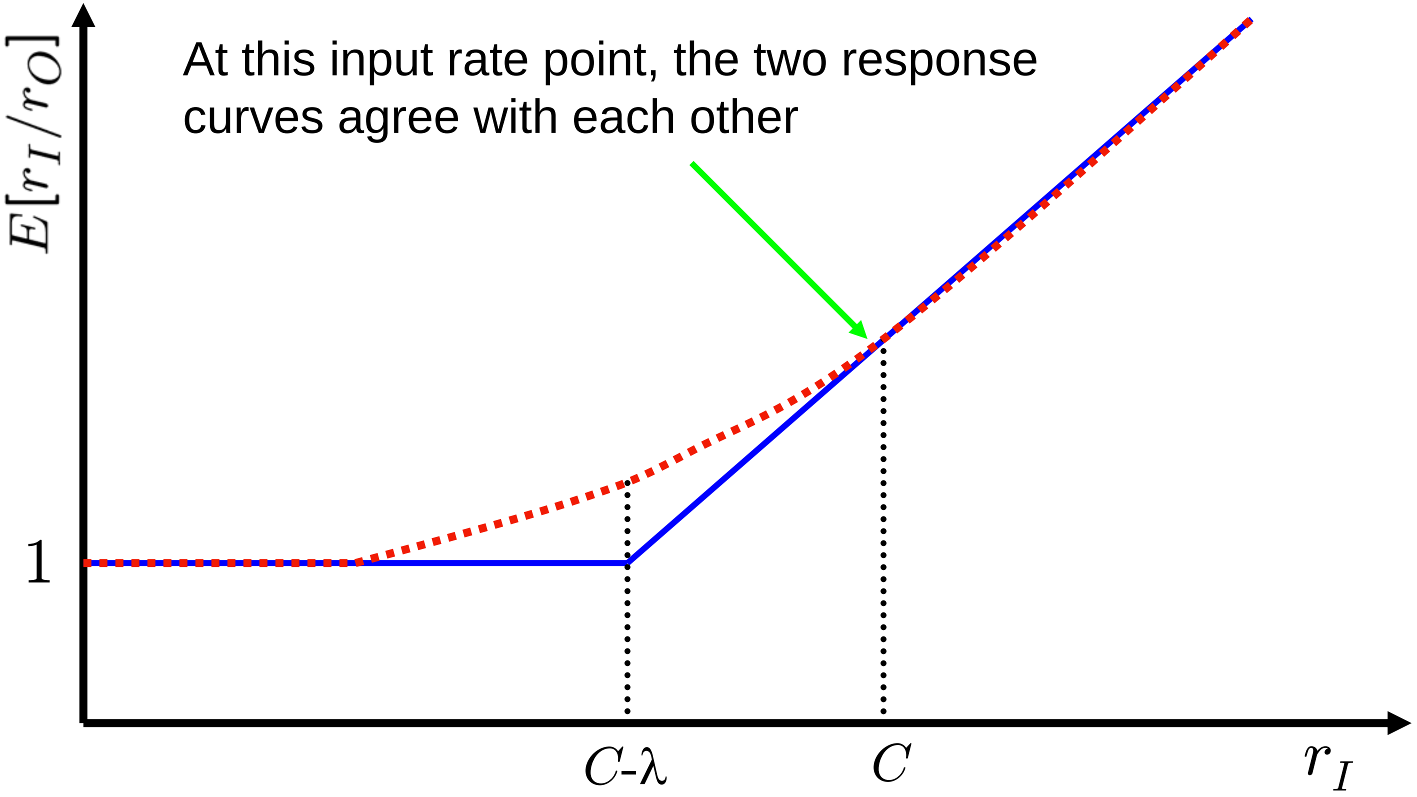

Implications: Spruce 2

- Measurement biases of Spruce in multi-hop paths.

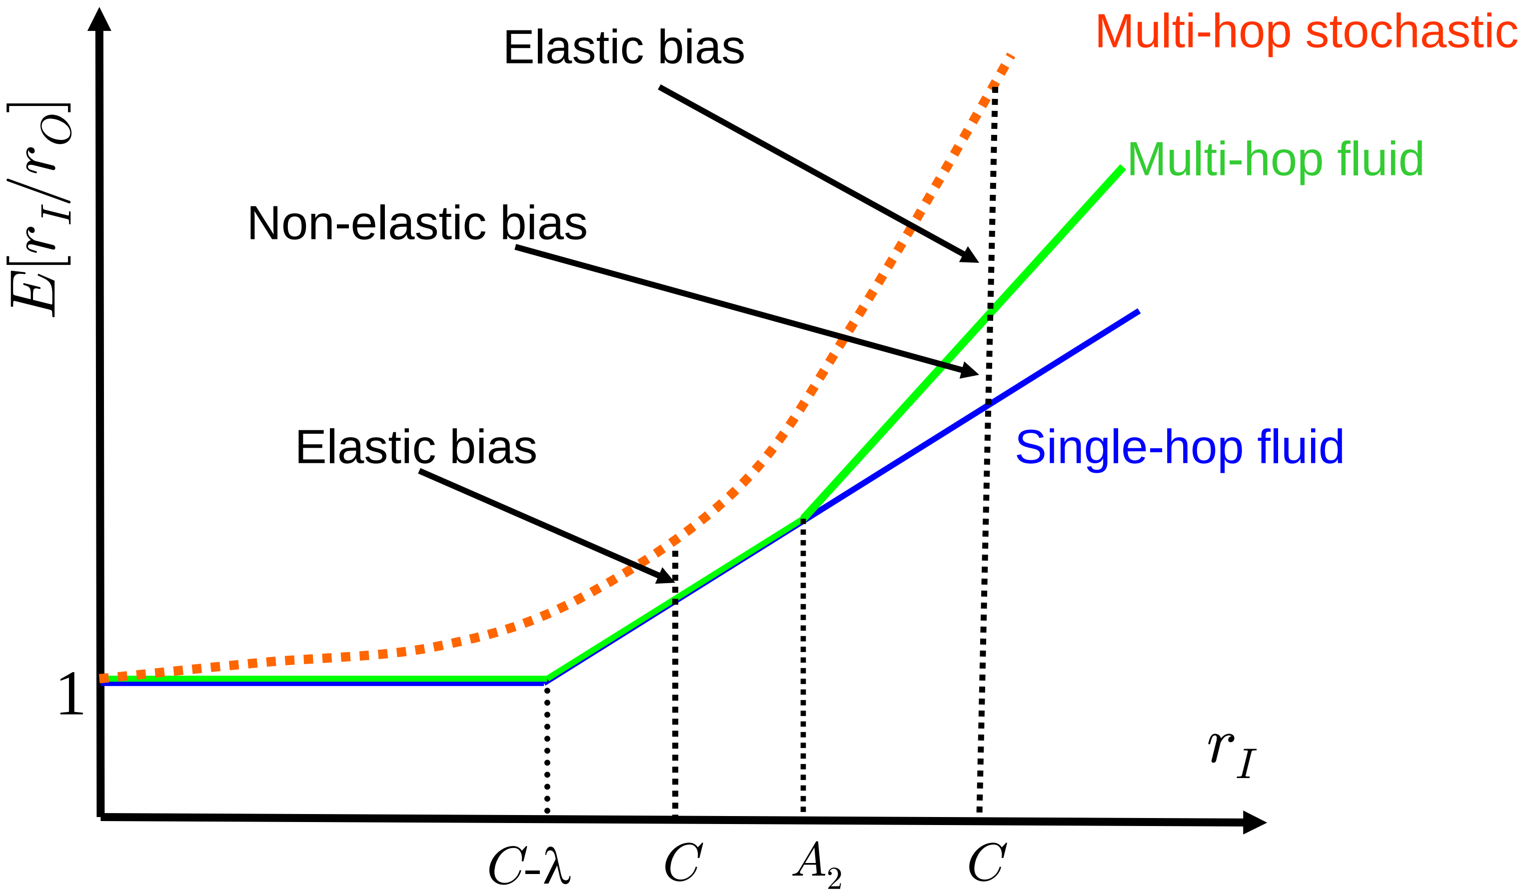

Implications: Spruce 3

- 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
Cornell-CMU	38	0	38	96	58

Implications: PTR and pathload

- 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, even in a multi-hop path.
  – Recall that using long trains, the multi-hop stochastic curve will approach the single-hop fluid curve within that region.

Wrap-up

- Conclusion
  - We derived the stochastic probing response curves in both single-hop and multi-hop paths
  - Our results provide a stochastic justification of the existing techniques using long-trains
  - Also uncover the sources of measurement biases for the techniques using short trains and possible ways to overcome the biases
  - Lead to new approaches for measuring the tight link capacity

Wrap-up

- Several Related Papers
  - X. Liu, K. Ravindran, B. Liu, and D. Loguinov, "Single-Hop Probing Asymptotics In Available Bandwidth Estimation: Sample-Path Analysis," *ACM IMC* 2004
  - X. Liu, K. Ravindran, and D. Loguinov, "Multi-Hop Probing Asympotics in Available Bandwidth Estimation: Stochastic Analysis," *ACM IMC* 2005
  - X. Liu, K. Ravindran, and D. Loguinov, "Measuring Probing Response Curves over the RON Testbed," *PAM* 2006