On Estimating Tight-Link Bandwidth Characteristics over Multi-Hop Path

Seong-Ryong Kang

Joint work with Xiliang Liu, Amit Bhati, and Dmitri Loguinov

Internet Research Lab
Department of Computer Science
Texas A&M University, College Station, TX 77843

July 6, 2006

Agenda

- Introduction
  - Motivations and goals
- Background
  - Definition of bandwidth
  - End-to-end Internet-path and single-hop model
- Envelope
  - Envelope packet-trains
  - Phase-based individual link measurement
- Performance of Envelope
  - Estimation accuracy and asymptotic behavior
- Wrap-up

Motivation

- Bandwidth estimation is an important area of Internet research
  - Helps to understand network path characteristics
  - Potentially can help various Internet applications

- Majority of available bandwidth estimation processes do not provably converge to the correct values

Motivation 2

- Furthermore, none of the existing techniques can correctly measure the tight-link capacity over a multi-hop path
  - Tight link is the link with the minimum available bandwidth of a path

Goals

- Develop a provably accurate estimation technique
- Measure both capacity and available bandwidth of the tight link
- Work for multi-hop paths under arbitrary cross-traffic

Definition of Bandwidth

- Bottleneck bandwidth
  - The capacity of the slowest link of an end-to-end path
  - The slowest link is often called *narrow* link
- Available bandwidth
  - The smallest average unused bandwidth along the end-to-end path

Definition Bandwidth 2

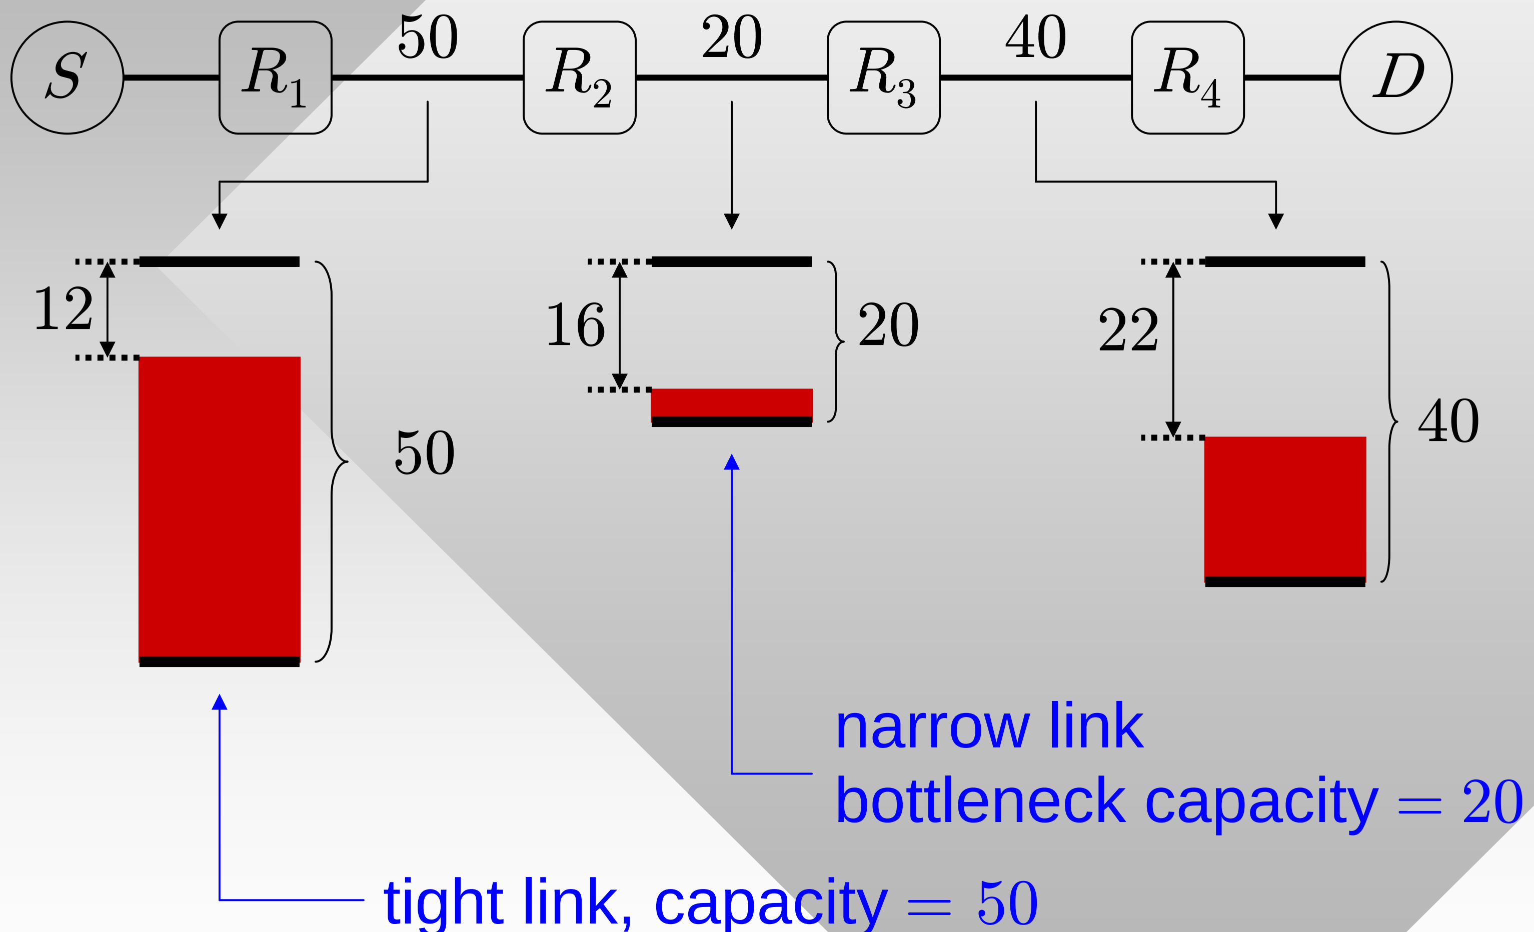

- Available bandwidth of the path: $A = 12$

End-to-End Internet Path

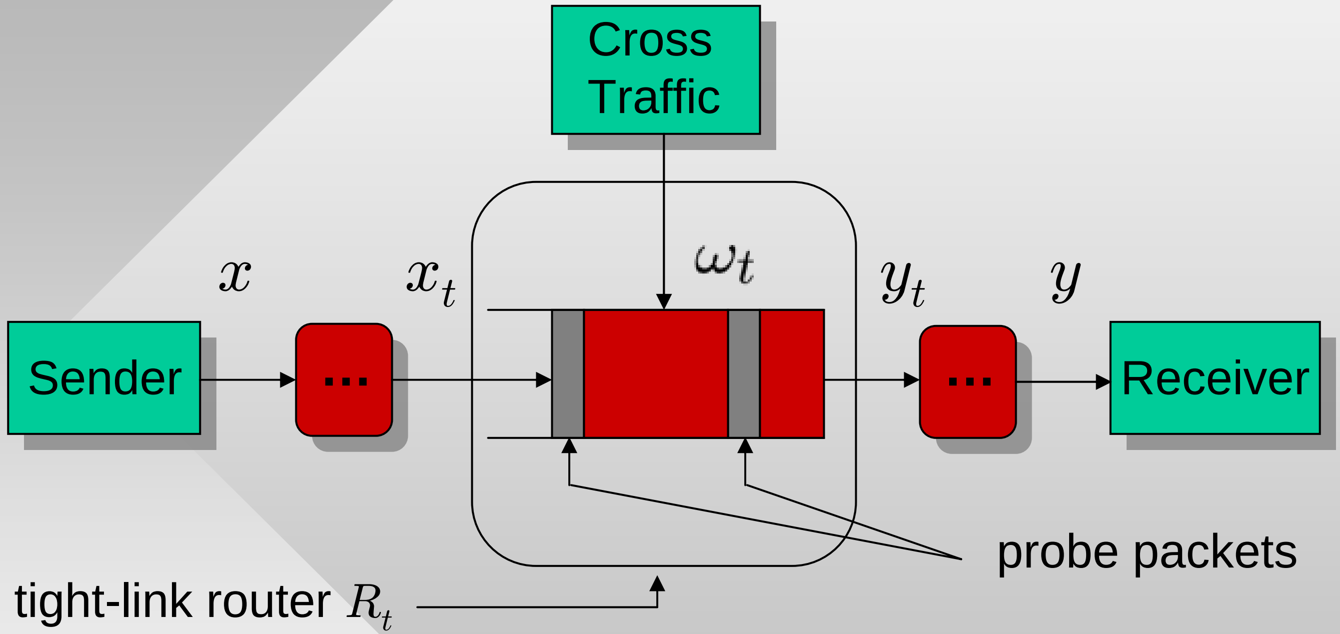

- The sender injects probe packets with inter-packet spacing x
- Due to expansion/compression in pre-tight links, inter-packet spacing x_t is different from the initial spacing x
- x_t is altered at router R_t to be y_t by random noise w_t
- The receiver samples inter-arrival dispersion y

Single-Hop Model

- Assumes that cross-traffic in non-tight links does not change inter-packet spacings of the probe packets
  - That is, $x = x_t$ and $y = y_t$
- Derives the mean output dispersion $E[y]$ under arbitrary cross-traffic
  - Other single-hop models (Dovrolis *et al*. INFOCOM 2001 and Melander *et al*. GLOBECOM 2000) rely on a constant-rate fluid model of cross-traffic

Single-Hop Model 2

- Extracts capacity and available bandwidth of the tight link from $E[y]$ with high accuracy
- For details, see the paper (Kang *et al*. INCP 2004)

Envelope

- Recursively extends the single-hop results to multi-hop paths
  - Sample statistics of each hop independently
- Measures both capacity and available bandwidth of the tight link under arbitrary cross-traffic
  - Can also measure non-tight link in certain path and cross-traffic conditions
- Provides asymptotic accuracy
  - Estimates converge to the true value after a sufficiently long measurement process

Envelope 2

- Recursive extension:
  - Treats inter-packet spacing x_k of probe traffic arriving at router R_k as the inter-departure delay y_{k-1} of the previous router R_{k-1}

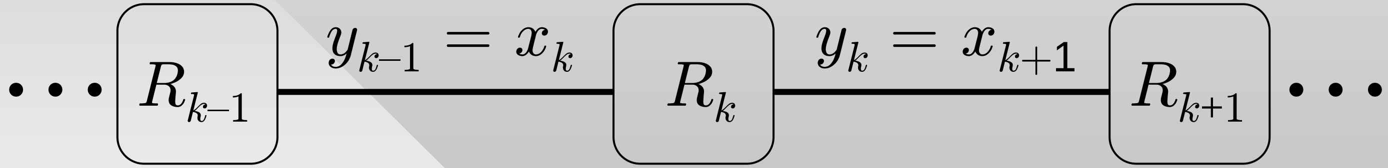

Envelope 3

- Necessary conditions for measuring link R_k
  - Spacing between two probe packets must be small when arriving at router R_k
  - However, the departure spacing from the router must be large to preserve its mean along the path suffix
- How do we satisfy these conditions?
  - By using Envelope packet-trains and TTL-limited dropping of probe packets at select routers
  - This does not require special “cooperation” from the routers

Envelope 4

- Envelope packet train

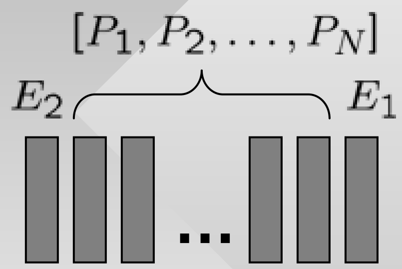

  - An envelope packet-train includes N probe packets P_1, …, P_N surrounded by two Envelope packets E_1 and E_2
  - Delays between two probe packets are small
  - Delays between two Envelope packets become large by selecting a large N

Envelope 5

- To obtain the departure spacing from router R_k, all probe packets P_1, …, P_N are dropped at router R_{k+1} using TTL limiting

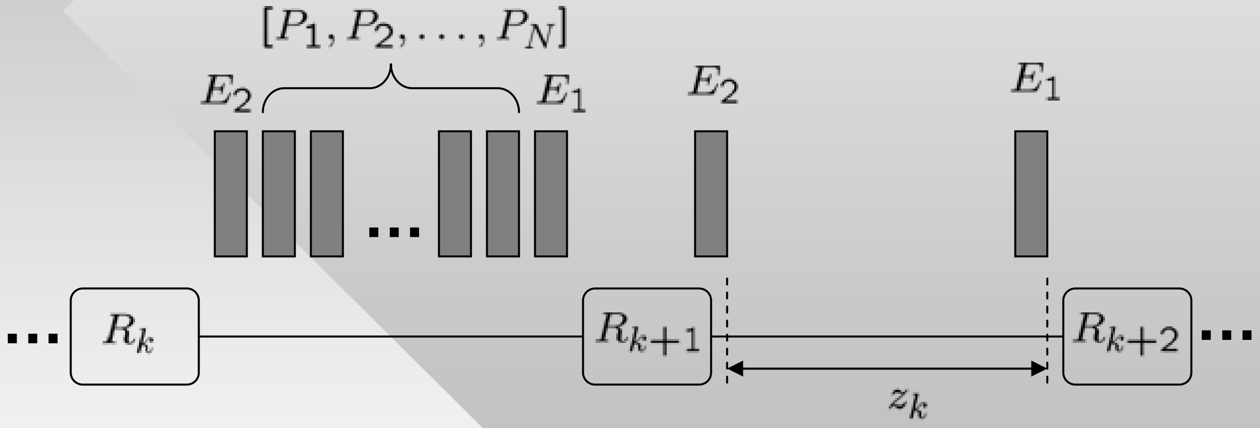

  - The probe packets sample queuing dynamics at the desired router R_k
  - The surviving envelope packets carry spacing z_k that is $N+1$ times larger than the departure spacing y_k from the router R_k

Envelope 6

- For M links in the end-to-end path, Envelope takes $M–1$ measurement phases
  - Each phase focuses on a particular router R_k
  - Obtains the mean spacing $E[z_k]$ exiting from R_k
- Using $E[z_{k-1}]$ measured in the previous phase, Envelope estimates the available bandwidth A_k and capacity C_k of the router R_k
- For details of bandwidth extraction process, see the paper

Simulation Topology

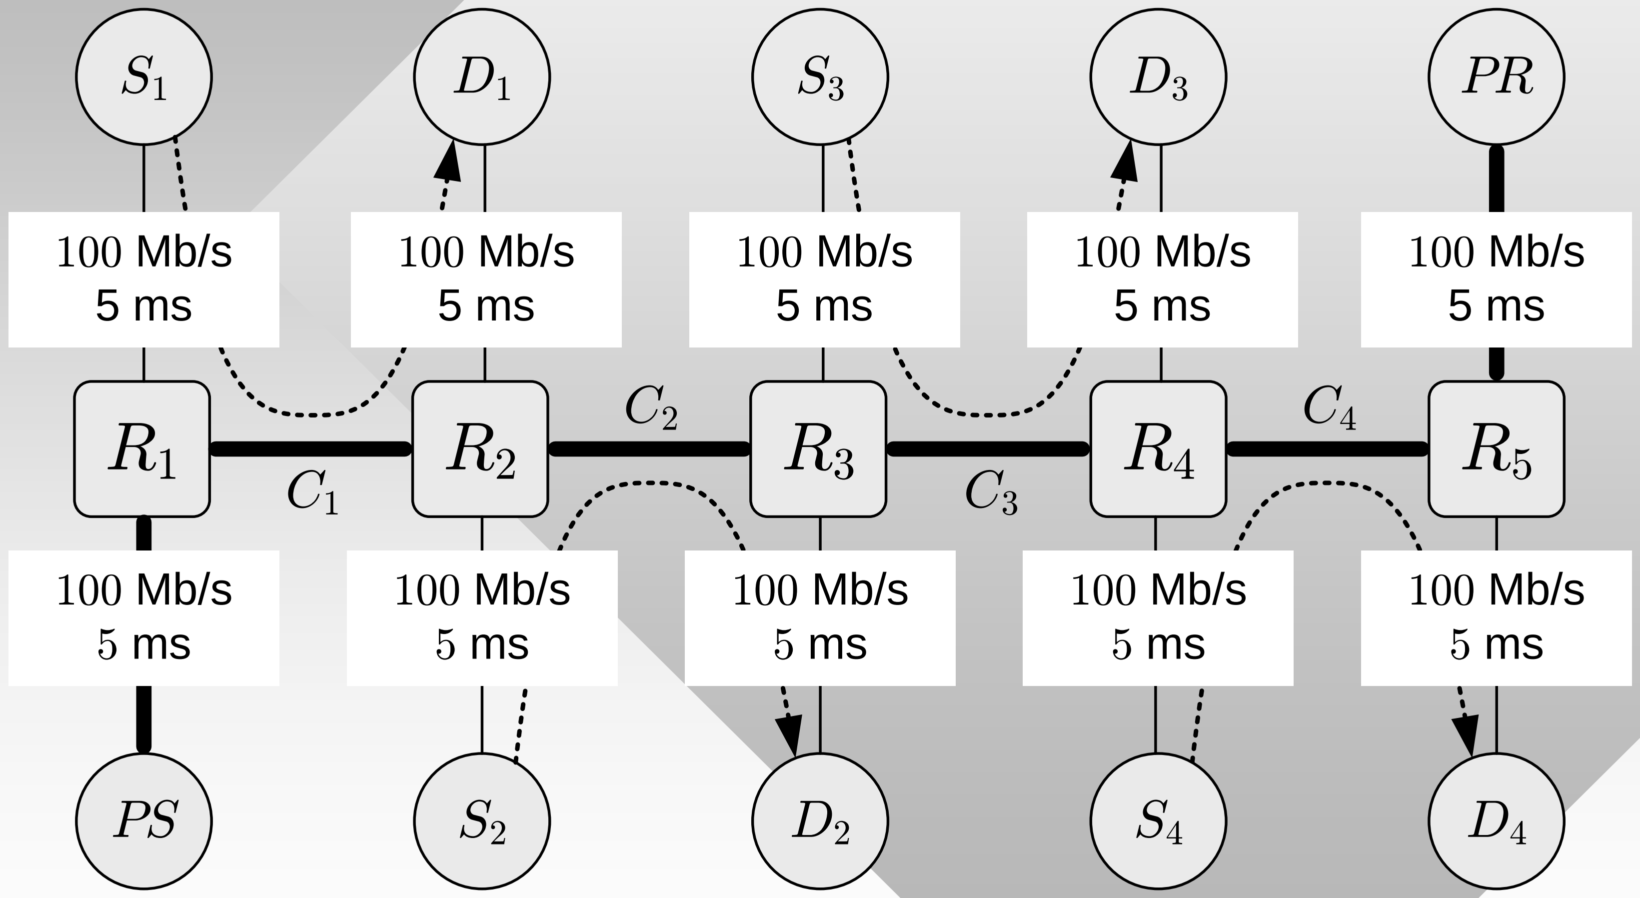

Simulation Setup

	Different link bandwidths (Mb/s)							
	C_1	A_1	C_2	A_2	C_3	A_3	C_4	A_4
Case-I	5	1	100	50	100	40	1.5	0.3
Case-II	2	0.4	1.5	0.25	0.8	0.4	1.5	0.35
Case-III	1.5	0.3	100	50	100	40	5	1
Case-IV	20	4	15	2.5	8	4	15	3.5
Case-V	2	0.4	0.8	0.4	1.5	0.25	2	0.4

– Darkly shaded values in each row represent the tight-link capacity and available bandwidth of the path

– We also lightly shade the narrow-link capacity in cases when it is different from the tight link

Estimation Accuracy

- Relative error metrics:

$$e_{C_i} = \frac{|C_i - C'_i|}{C_i}, \quad e_{A_i} = \frac{|A_i - A'_i|}{A_i}$$

  - e_{C_i} and e_{A_i} are relative estimation errors of C_i and A_i, respectively
  - C_i is the true capacity of link i and C_i' is its estimate
  - A_i is the true available bandwidth of link i and A_i' is its estimate

Estimation Accuracy 2

- CBR cross-traffic

	Relative estimation error				
	Case-I	Case-II	Case-III	Case-IV	Case-V
e_{C_1}	0.94%	2.39%	0.17%	0.15%	10.76%
e_{A_1}	7.75%	1.57%	3.74%	6.99%	4.20%
e_{C_2}	—	0.35%	—	2.36%	2.47%
e_{A_2}	—	2.09%	—	5.62%	8.71%
e_{C_3}	—	3.76%	—	0.65%	4.13%
e_{A_3}	—	7.07%	—	2.04%	5.71%
e_{C_4}	1.56%	0.60%	—	12.11%	21.19%
e_{A_4}	2.38%	3.05%	—	9.86%	17.59%

Estimation Accuracy 3

- TCP cross-traffic

	Relative estimation error				
	Case-I	Case-II	Case-III	Case-IV	Case-V
e_{C_1}	0.21%	0.40%	0.46%	8.55%	4.22%
e_{A_1}	12.07%	0.27%	0.98%	21.23%	16.60%
e_{C_2}	—	3.62%	—	0.26%	5.90%
e_{A_2}	—	4.22%	—	3.29%	10.06%
e_{C_3}	—	10.79%	—	9.41%	10.06%
e_{A_3}	—	15.44%	—	23.30%	5.82%
e_{C_4}	0.24%	10.04%	—	—	—
e_{A_4}	3.30%	11.53%	—	—	—

Asymptotic Behavior

- CBR cross-traffic (case I)

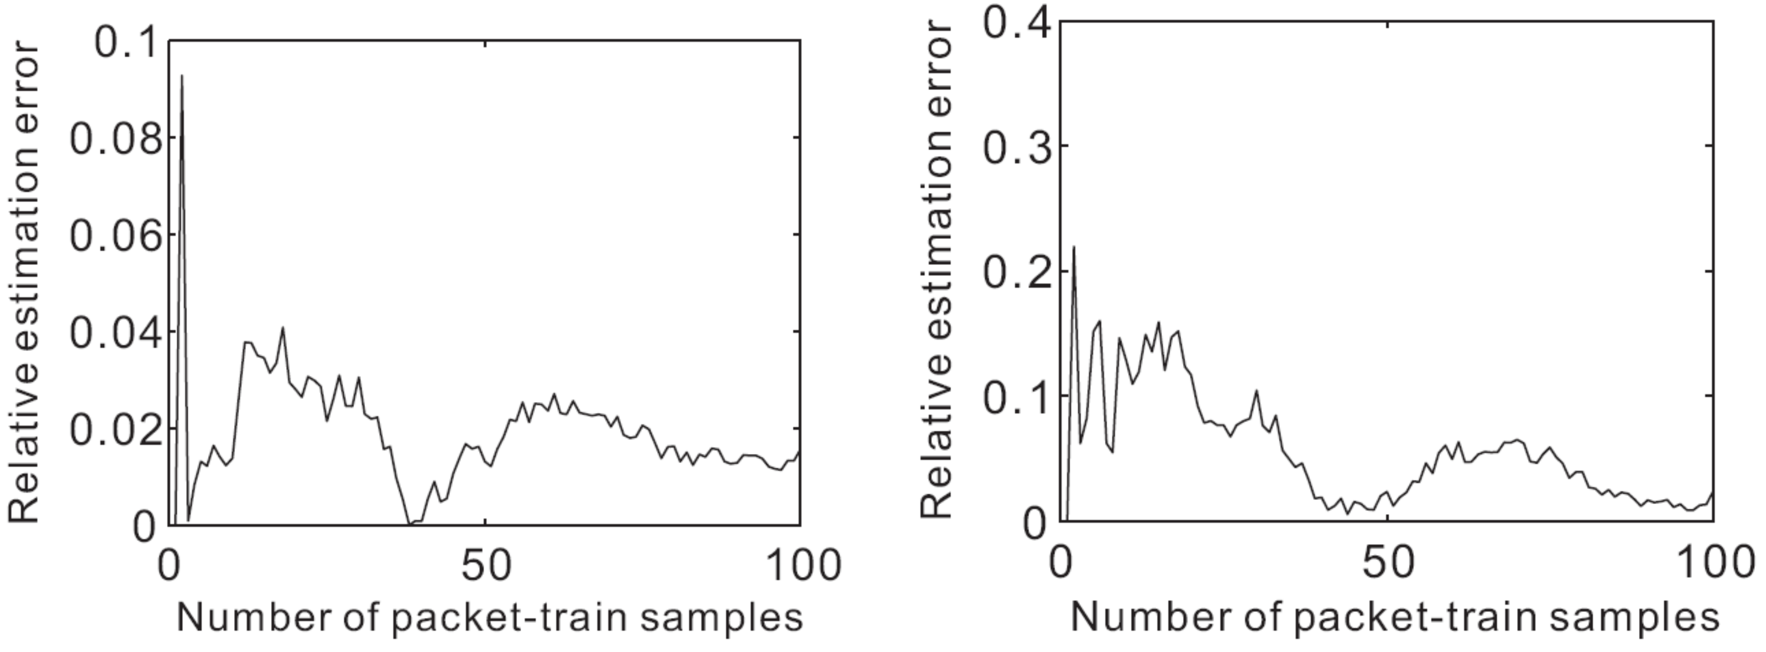

(a) e_C for case I (b) e_A for case I

Asymptotic Behavior 2

- TCP cross-traffic (case I)

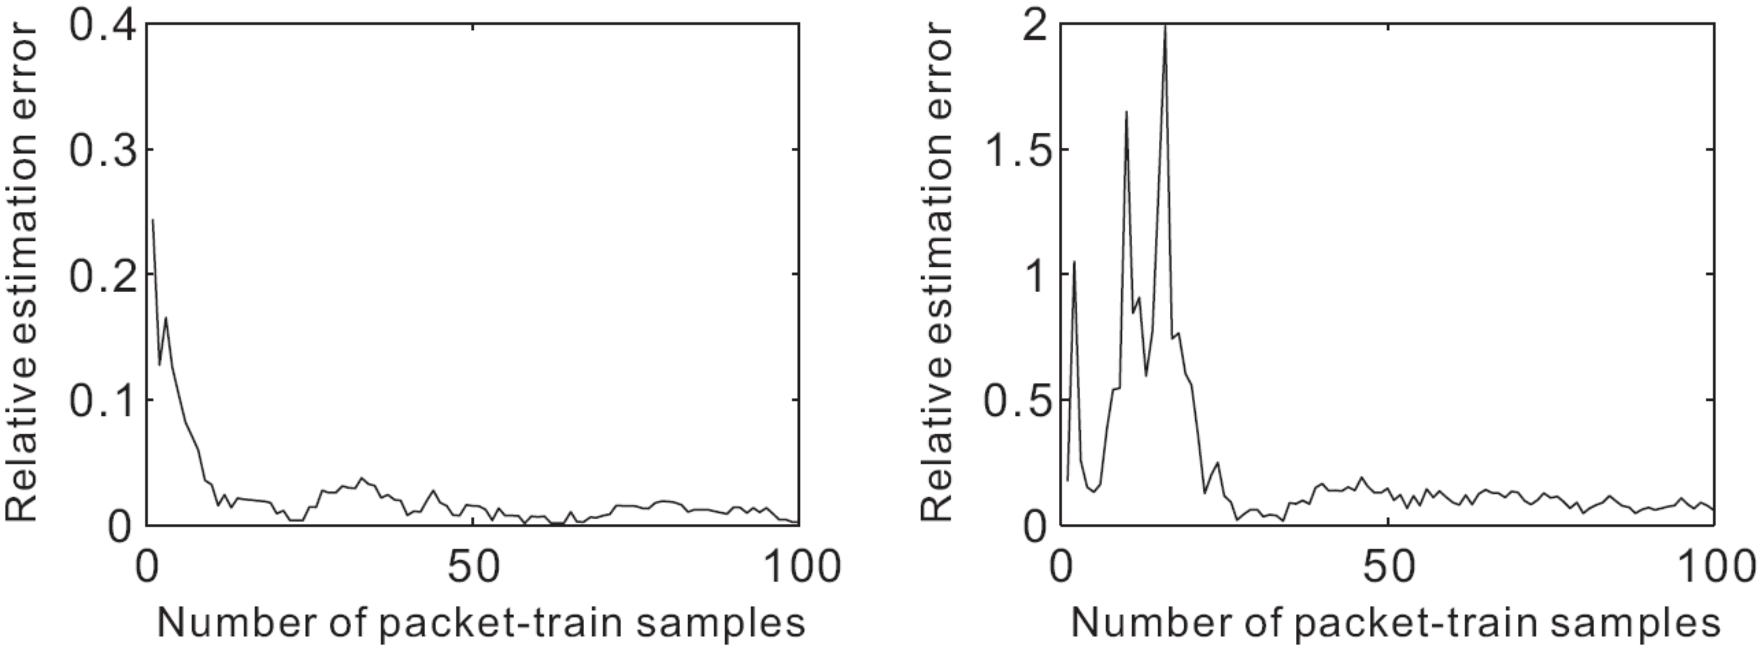

(a) e_C for case I

(b) e_A for case I

Performance Comparison

- Available bandwidth under TCP cross-traffic

	Relative estimation error			
	Envelope	Pathload	Spruce	IGI
Case-I	3.30%	8.33%	34.83%	86.67%
Case-II	4.22%	12.09%	78.00%	109.20%
Case-III	0.98%	3.33%	7.65%	103.05%
Case-IV	3.29%	15.88%	78.15%	98.96%
Case-V	5.82%	12.04%	70.85%	91.60%

Performance Comparison 2

- Bottleneck bandwidth under TCP cross-traffic

	Relative estimation error		
	Envelope	CapProbe	Pathrate
Case-I	0.24%	40.95%	40.93%
Case-II	10.79%	39.12%	32.50%
Case-III	0.46%	35.78%	48.10%
Case-IV	9.41%	50.60%	20.62%
Case-V	5.90%	51.62%	45.62%

Wrap-up

- Envelope measures both bandwidth metrics of the tight-link under arbitrary cross-traffic
- Its estimates are asymptotically accurate
- It is based on recursive extension of the single-hop results to multi-hop paths
- Future work
  - Implementation and deployment of Envelope
  - Further reduction of probe traffic required