End-to-end available bandwidth estimation

Constantinos Dovrolis

Computer and Information Sciences

University of Delaware

The big picture

- The Internet is really a big **black** cloud

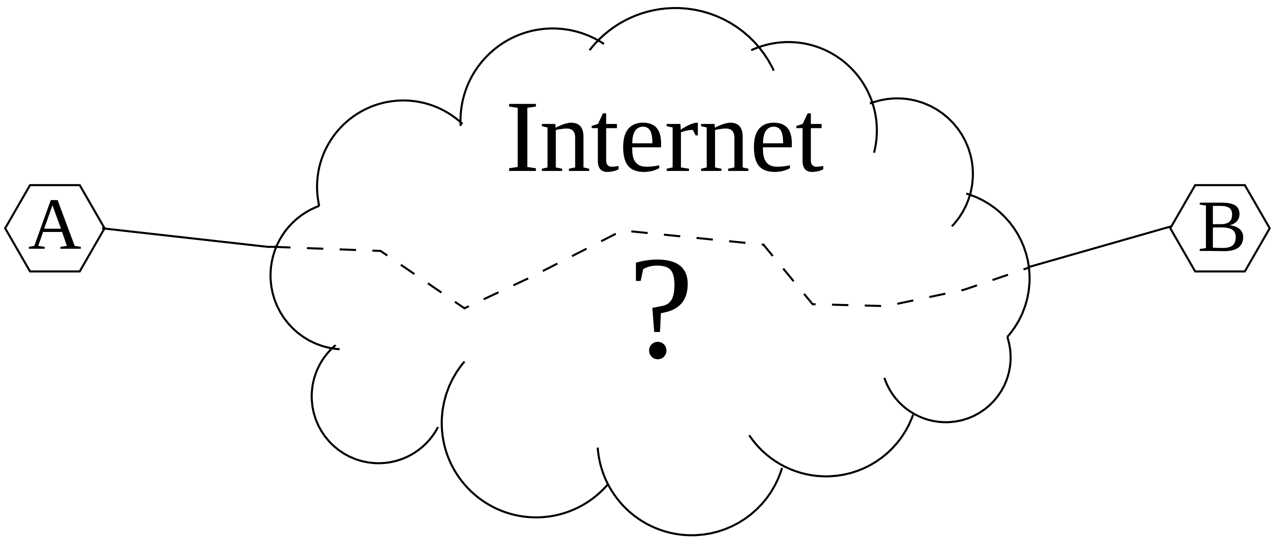

But..

- End-systems can *infer* what is in the box through end-to-end measurements
- Why bother? Improved transport, QoS, overlay routing, server selection, etc
- This work focuses on *end-to-end bandwidth estimation*

Overview

- Bandwidth metrics: capacity and available bandwidth
- Available bandwidth measurement methodology: *SLoPS*
- Available bandwidth measurement tool: *pathload*
- Using *pathload* to understand the dynamics of available bandwidth

Part I

Capacity and Available Bandwidth

Capacity

- **Capacity:** maximum possible end-to-end throughput

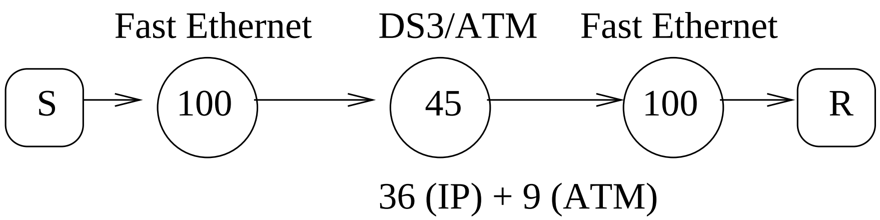

- End-to-end capacity C is limited by *narrow link* n:

$$C = \min_{i=0 \ldots H} \{C_i\} = C_n$$

- **Pathrate:** measurement tool based on packet pairs/trains (Infocom'01)

  See www.pathrate.org

Available bandwidth

- **Avail-bw:** maximum end-to-end throughput without reducing cross traffic rate

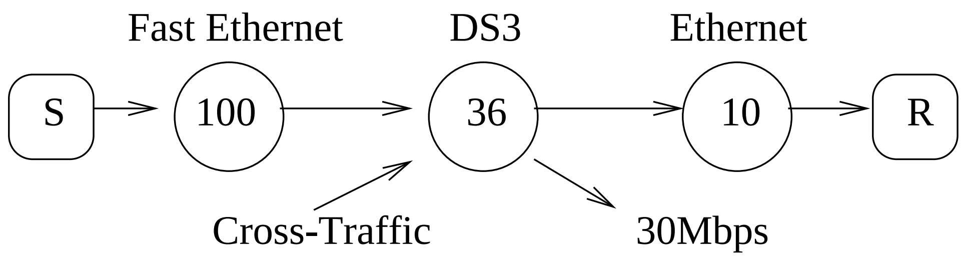

- Network managers monitor avail-bw with MRTG (based on router statistics)

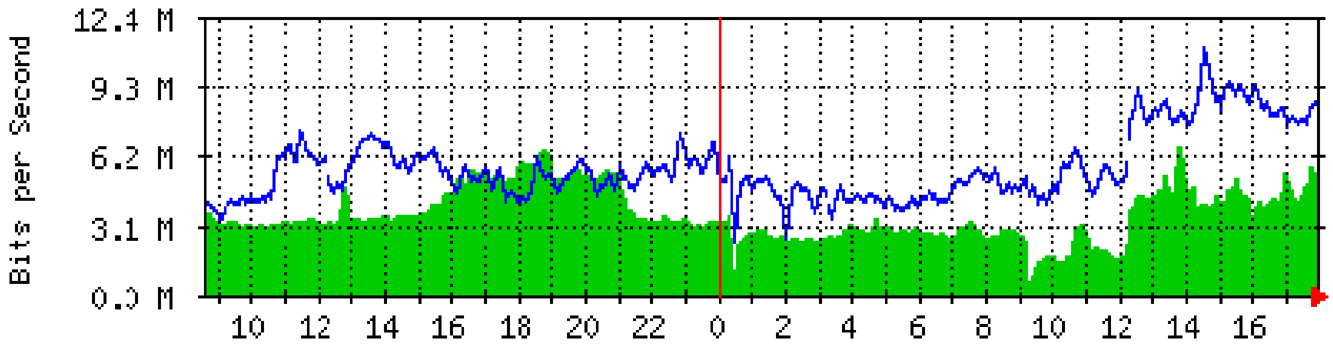

Definition of avail-bw

- C_i: capacity of link i
- u_i: utilization of link i in time interval T $(0 \leq u_i \leq 1)$
- Avail-bw of link i during T: $A_i = C_i\,(1 - u_i)$

$$A = \min_{i=0 \ldots H} A_i = \min_{i=0 \ldots H} C_i\,(1 - u_i)$$

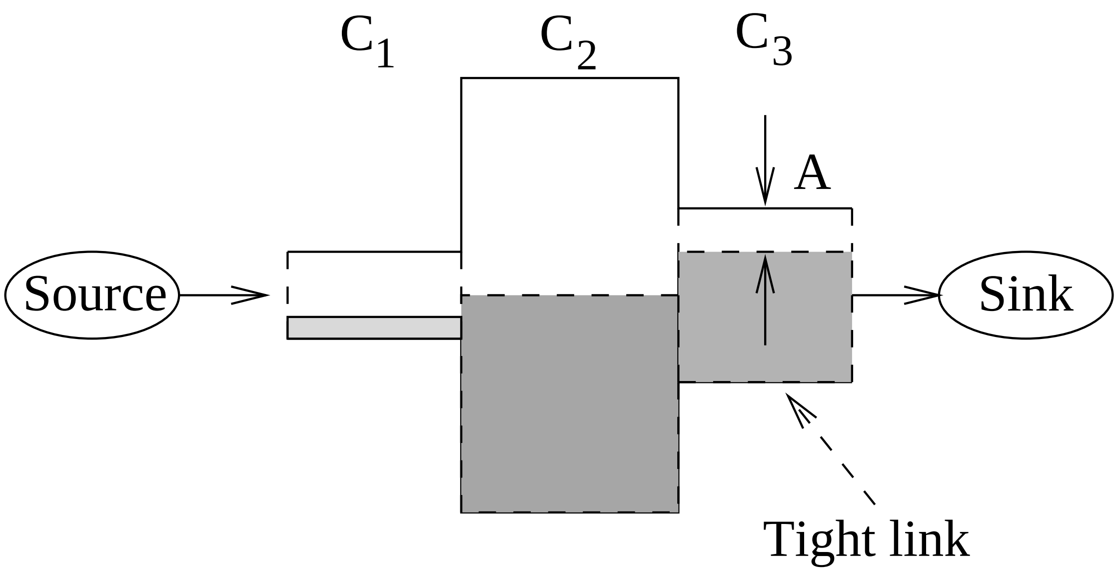

- Avail-bw is limited by *tight link*

Part II

Available Bandwidth Estimation

Applications of avail-bw estimation

- Congestion control and TCP: measure *Bandwidth-Delay-Product*
- Streaming applications: adjust encoding rate
- SLA and QoS verification: monitor path load
- Content distribution networks: select best server
- Overlay networks: configure overlay routes
- End-to-end admission control: check for sufficient bandwidth
- But how can we measure end-to-end avail-bw?

Previous work on available bandwidth estimation

1. Measure throughput of large TCP transfer
   - TCP does not get available bandwidth in under-buffered paths
   - TCP gets more than available bandwidth in over-buffered paths
   - TCP saturates the path *(intrusive measurements)*
2. Carter & Crovella: dispersion of long packet trains (cprobe)

   Does not measure available bandwidth (see Infocom'01)
3. Banerjee & Agrawala (ICN'00): define *available capacity* as:

   Amount of data that can be sent to path with OWD of less than Δ
4. Melander et.al. (Global Internet'00) and Ribeiro et.al. (ITC'00)

   Correct estimation when queueing only at single link in path

Does bulk-TCP measure avail-bw?

- Conventional wisdom: bulk-TCP oscillates around avail-bw
- Perform bulk-TCP transfer during (B) and (D) (5-min intervals)

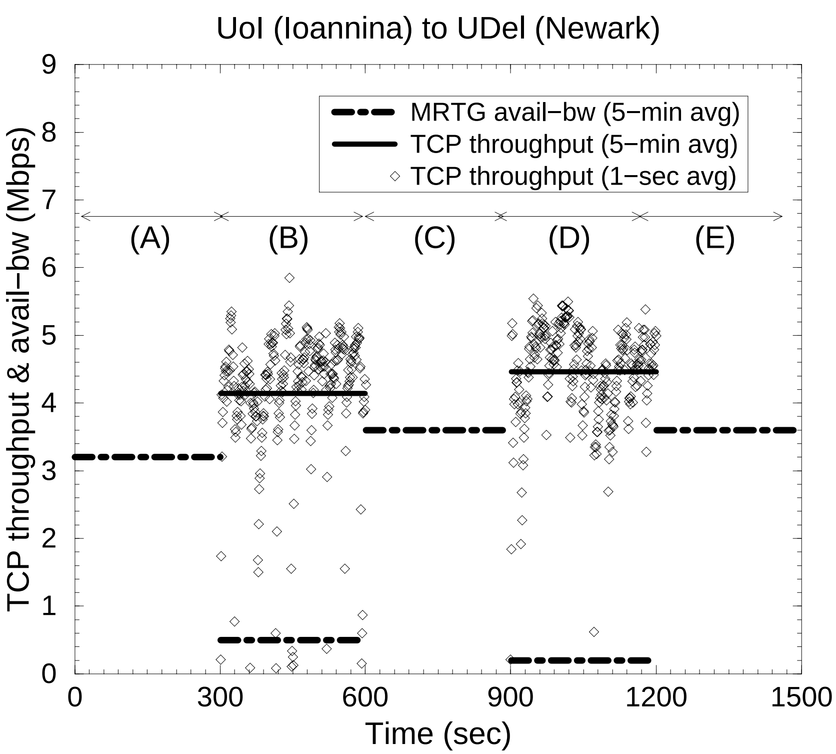

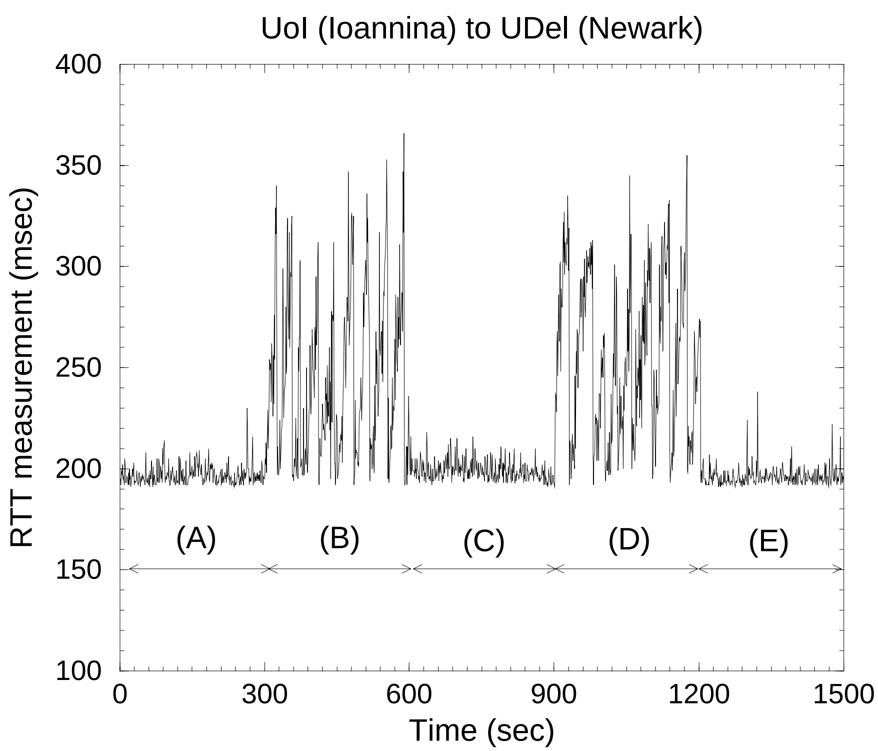

Self-Loading Periodic Streams (SLoPS)

- SLoPS requires access at both ends S and R of path
- S sends periodic UDP packet streams to R (timestamped)
- SLoPS analyzes One-Way Delays (OWDs) of packets from S to R
- OWD: $D_i = T^R_{arrive} - T^S_{send} = T_{arrive} - T_{send} + \text{Clock\_Offset}(S, R)$
- Interested in OWD variations: $D_i - D_{i+1}$
- So, S and R do NOT need to have synchronized clocks

SLoPS: Basic idea

- Periodic stream: K packets, size L bytes, rate $R = L/T$

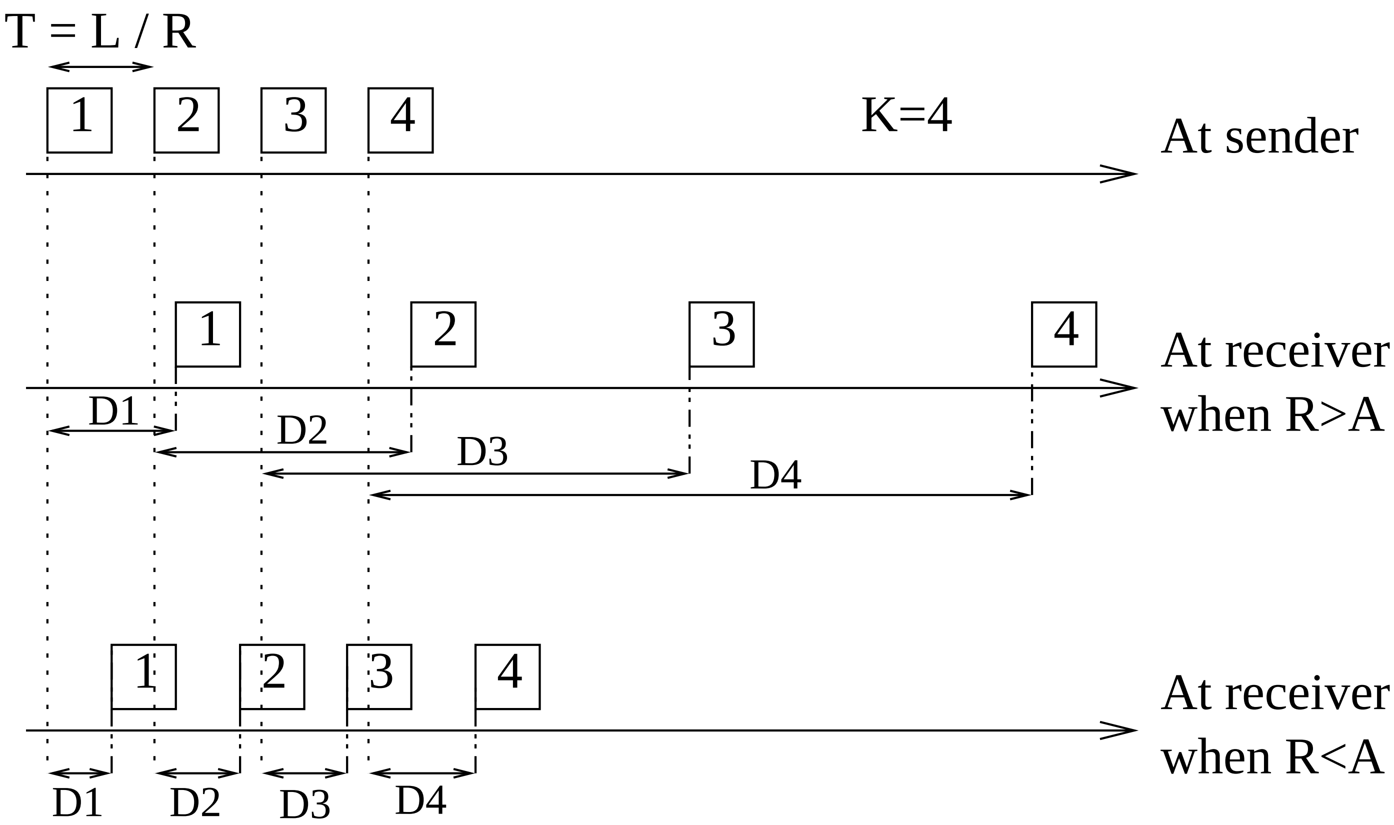

- **If $R > A$, OWDs gradually increase due to self-loading of stream**

Analytical model

- Flow conservation at tight link in interval $(t_i, t_{i+1} = t_i + T)$ with $T = \frac{L}{R}$:

$$\text{Arrivals} + \text{Queue} = \text{Departures} + \text{Queue}'$$

$$V(t_i, t_i + T) + Q(t_i) = S(t_i, t_i + T) + Q(t_i + T)$$

- Expected arrivals:
  $E[V(t_i, t_i + T)] = u_t C_t T + L = (C_t - A)T + RT = [C_t + (R - A)]T$
- Upper bound on departures: $S(t_i, t_i + T) \leq C_t T$
- If $R > A$,

$$E[V(t_i, t_i + T)] > S(t_i, t_i + T) \;\rightarrow\; Q(t_i + T) > Q(t_i)$$

- But, $Q(t_i + T) > Q(t_i) \;\rightarrow\; D_{i+1} > D_i$

  i.e., Packet stream has increasing OWDs

Increasing trend: $R > A$

- A=74Mbps, R=96Mbps, ($K = 100$ packets, T=100μs, L=1200B)

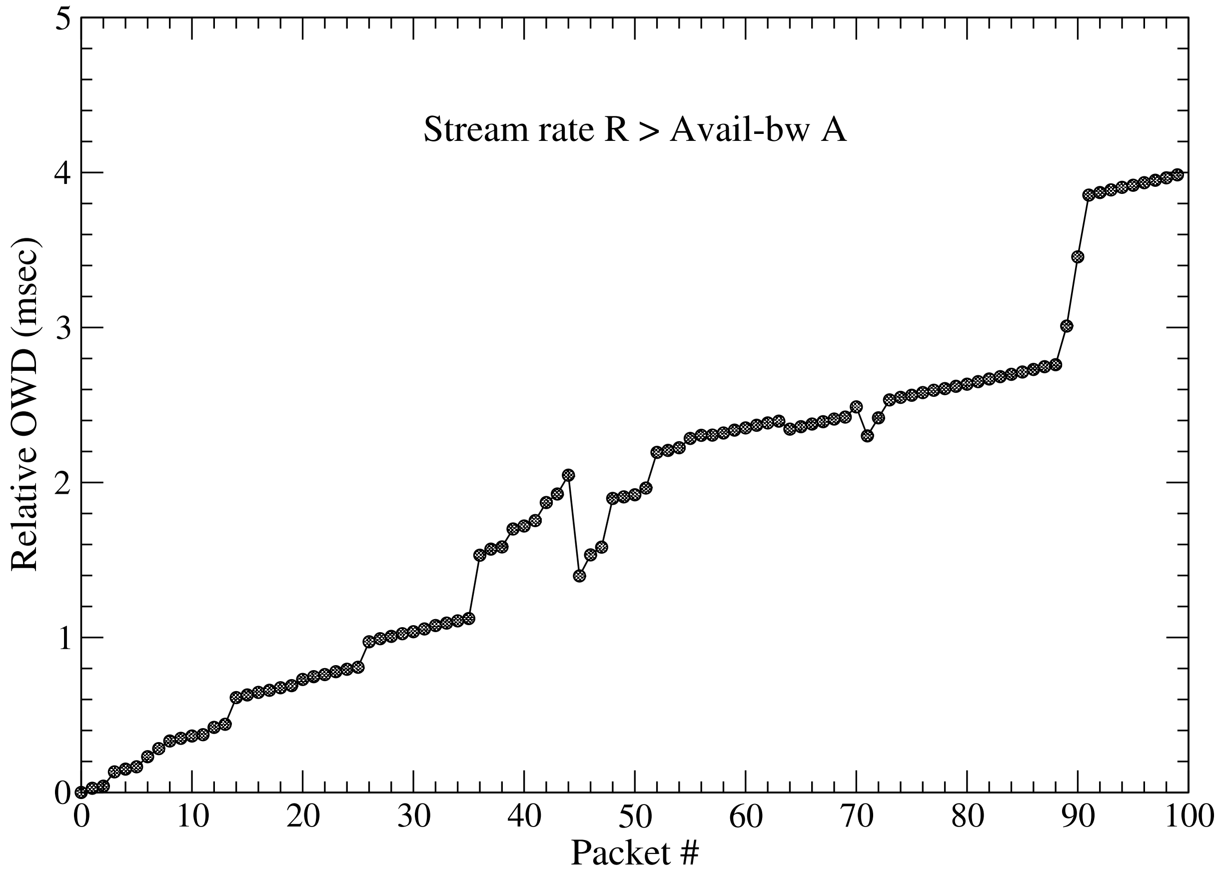

Non-increasing trend: $R < A$

- A=74Mbps, R=37Mbps, ($K = 100$ packets, T=100μs, L=462B)

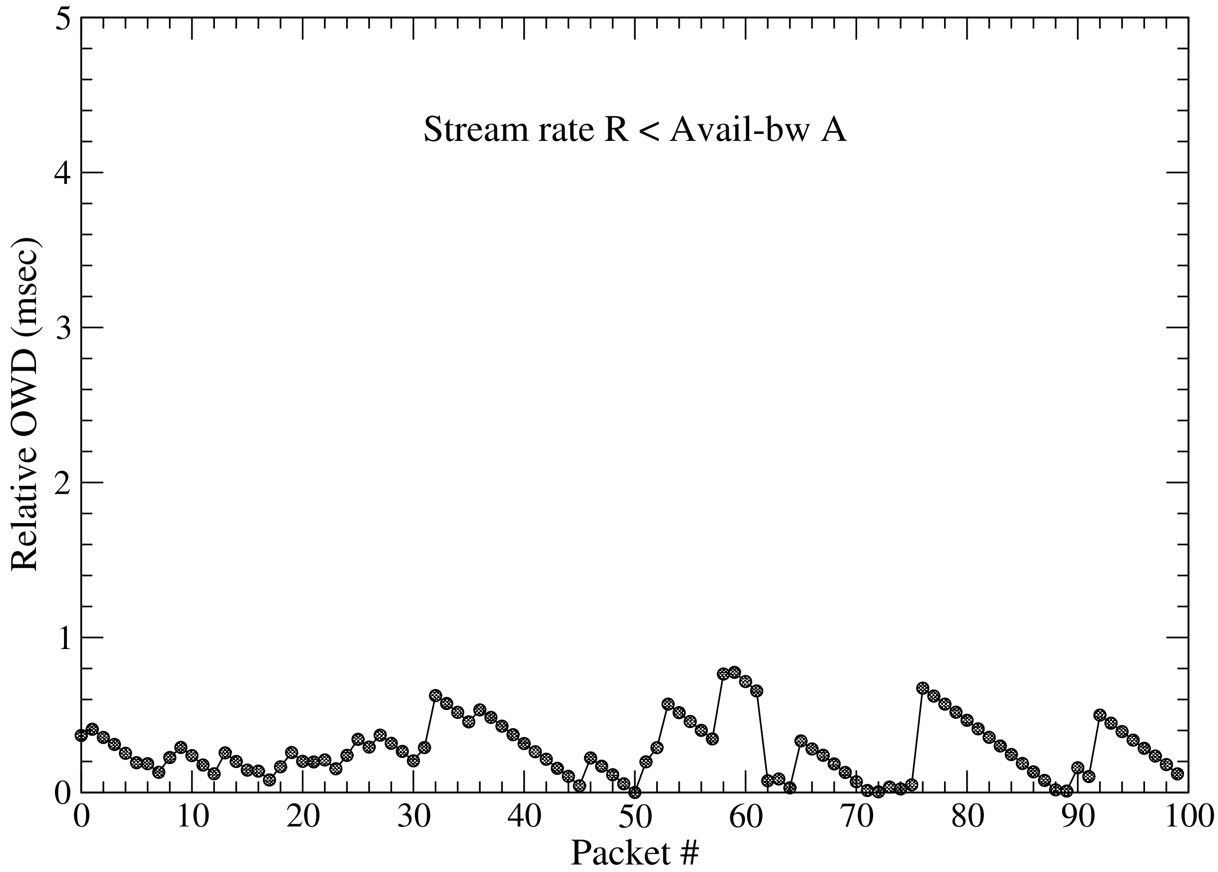

Grey-region: $R \bowtie A$

- A=74Mbps, R=82Mbps, ($K = 100$ packets, T=100μs, L=1025B)

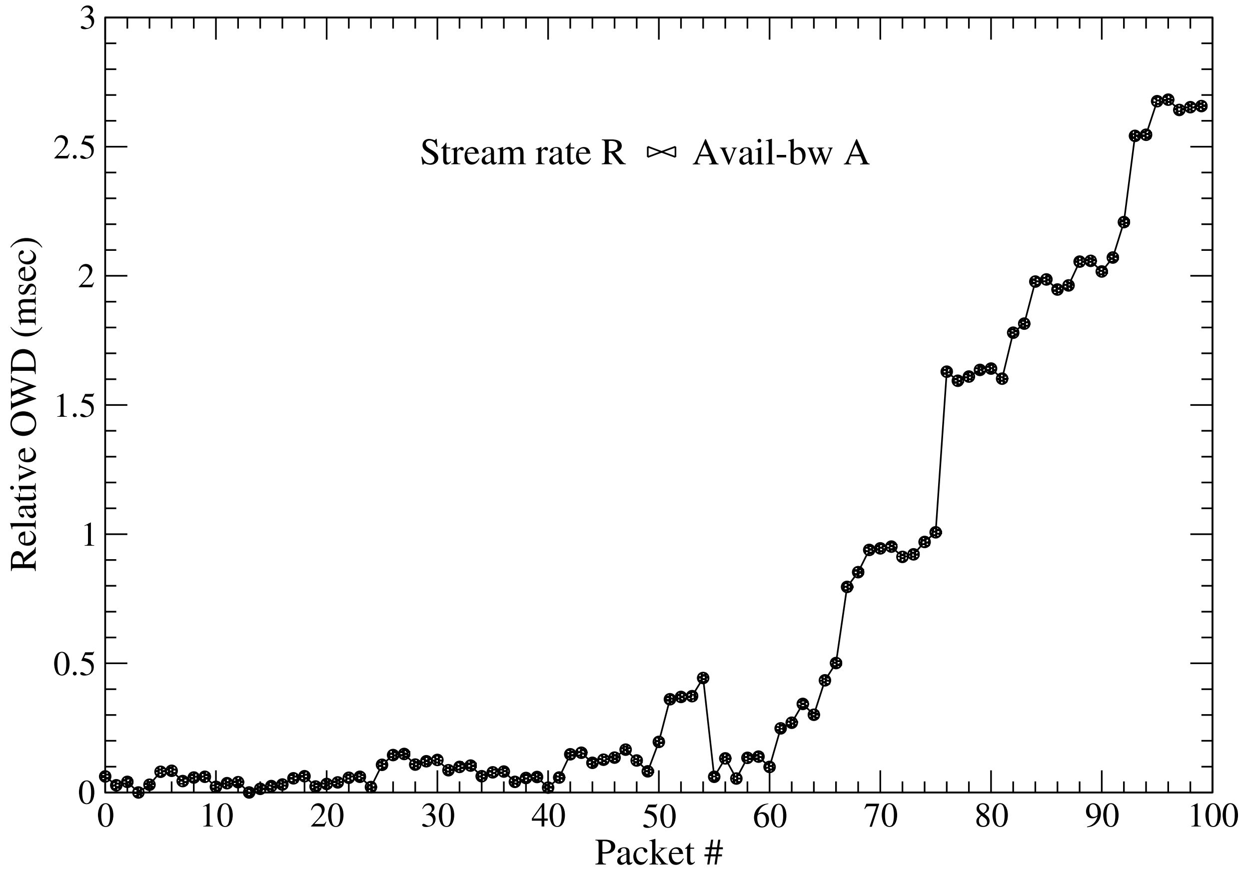

Part III

Pathload

Iterative algorithm in SLoPS

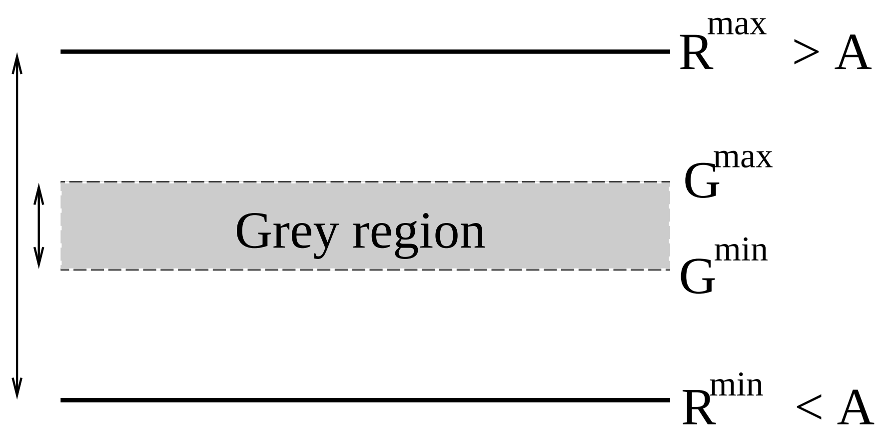

Terminate if:

$$R^{max} - R^{min} < w$$

or

$$G^{max} - G^{min} \cong R^{max} - R^{min}$$

Increasing trend: $R(n) > A$

$$R^{max} = R(n)$$

$$R(n+1) = (G^{max} + R^{max})/2$$

Non-increasing trend: $R(n) < A$

$$R^{min} = R(n)$$

$$R(n+1) = (G^{min} + R^{min})/2$$

Grey region & $R(n) > G^{max}$:

$$G^{max} = R(n)$$

$$R(n+1) = (G^{max} + R^{max})/2$$

Grey region & $R(n) < G^{min}$:

$$G^{min} = R(n)$$

$$R(n+1) = (G^{min} + R^{min})/2$$

Pathload features not covered in this talk

- Statistical tests for detection of increasing trend
- Clock-skew compensation
- Detection of context switches at sender & receiver
- Fleets: a number of streams of same rate (spaced by one RTT)
- Selection of packet size L and period T
- Dealing with losses and congestion responsiveness
- Initialization of rate adjustment algorithm

- For more details, see PAM 2002 publication on *pathload*

Verification results

- Verify avail-bw using MRTG graphs for path routers

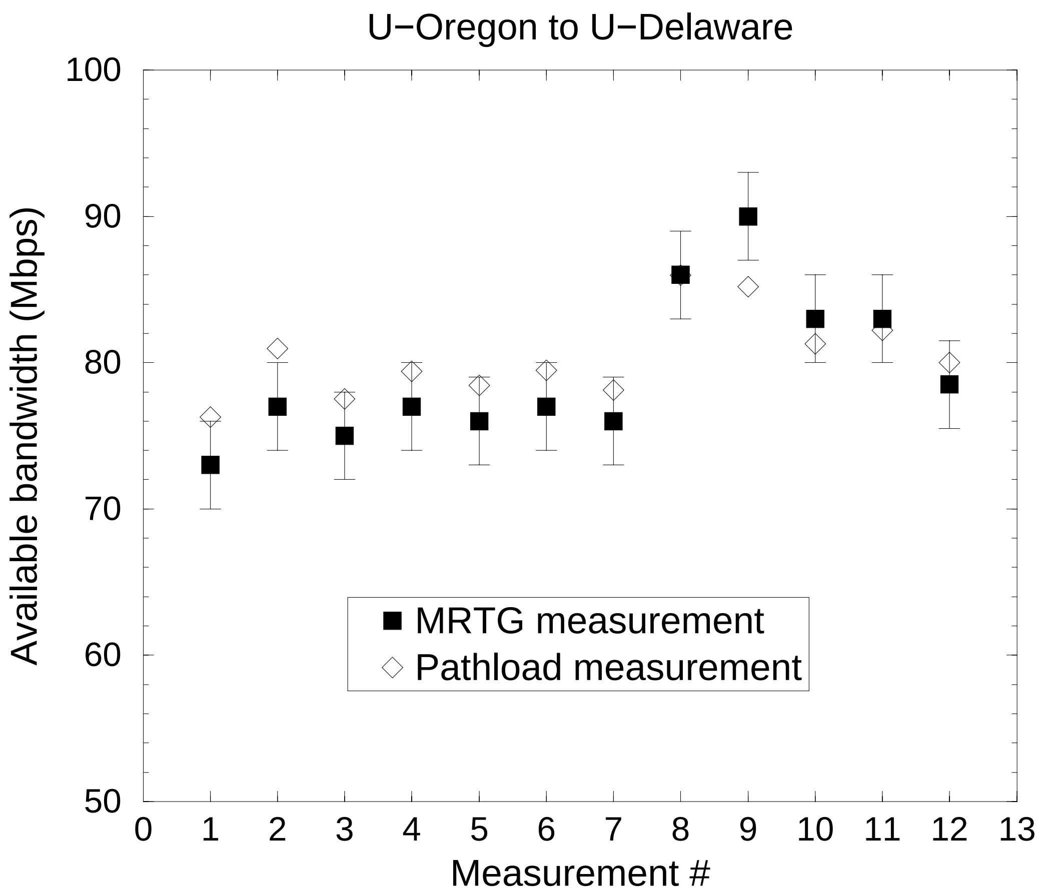

- Tight link: U-Oregon GigaPoP link (C=155Mbps), ω=3Mbps

Is pathload intrusive?

- Does pathload decrease avail-bw? Does it increase delays & losses?
- Run pathload during (B) and (D) (5-min intervals)

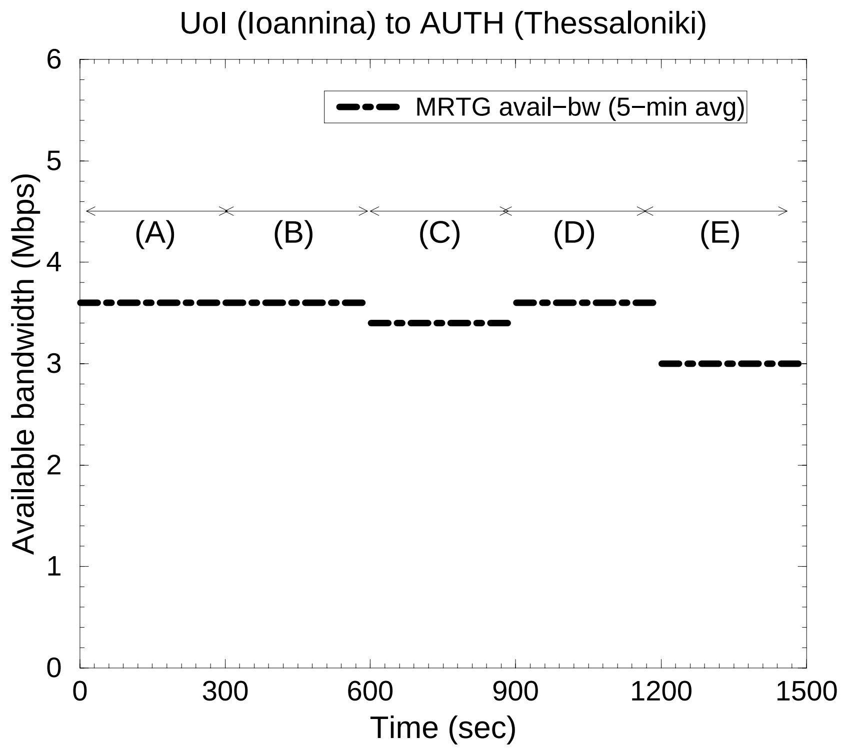

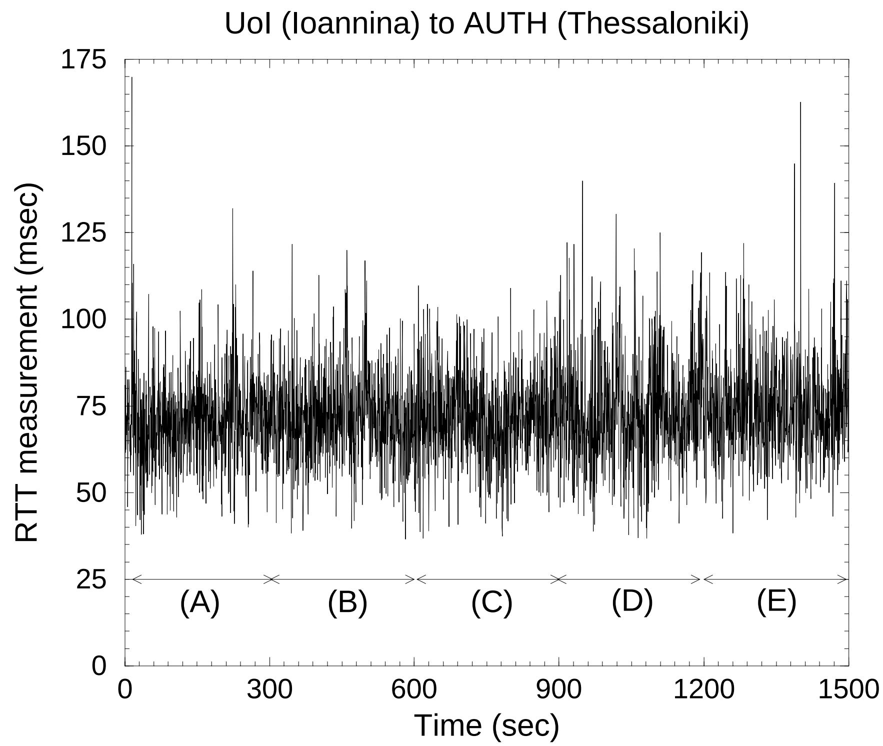

Part IV

Variability of available bandwidth

Variability and load conditions

- Relative variation of avail-bw: $\rho = \frac{R^{max} - R^{min}}{(R^{max} + R^{min})/2}$

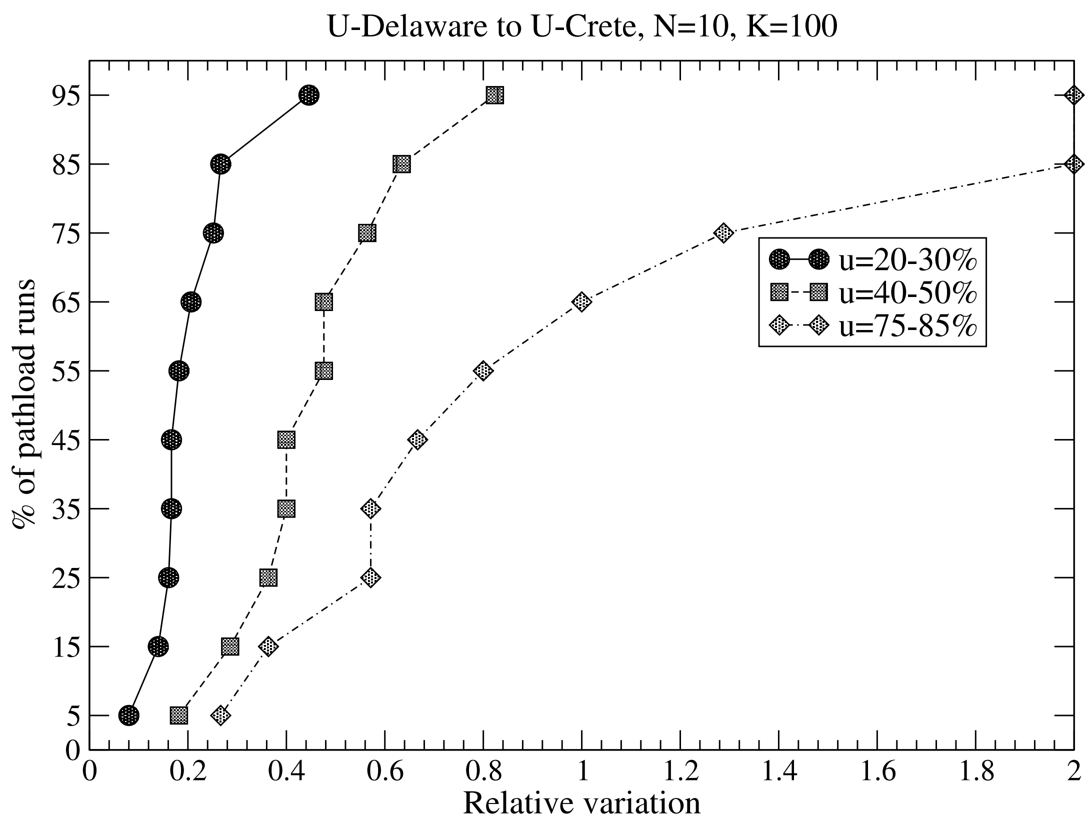

- Heavier utilization at tight link causes higher variability

Variability and statistical multiplexing

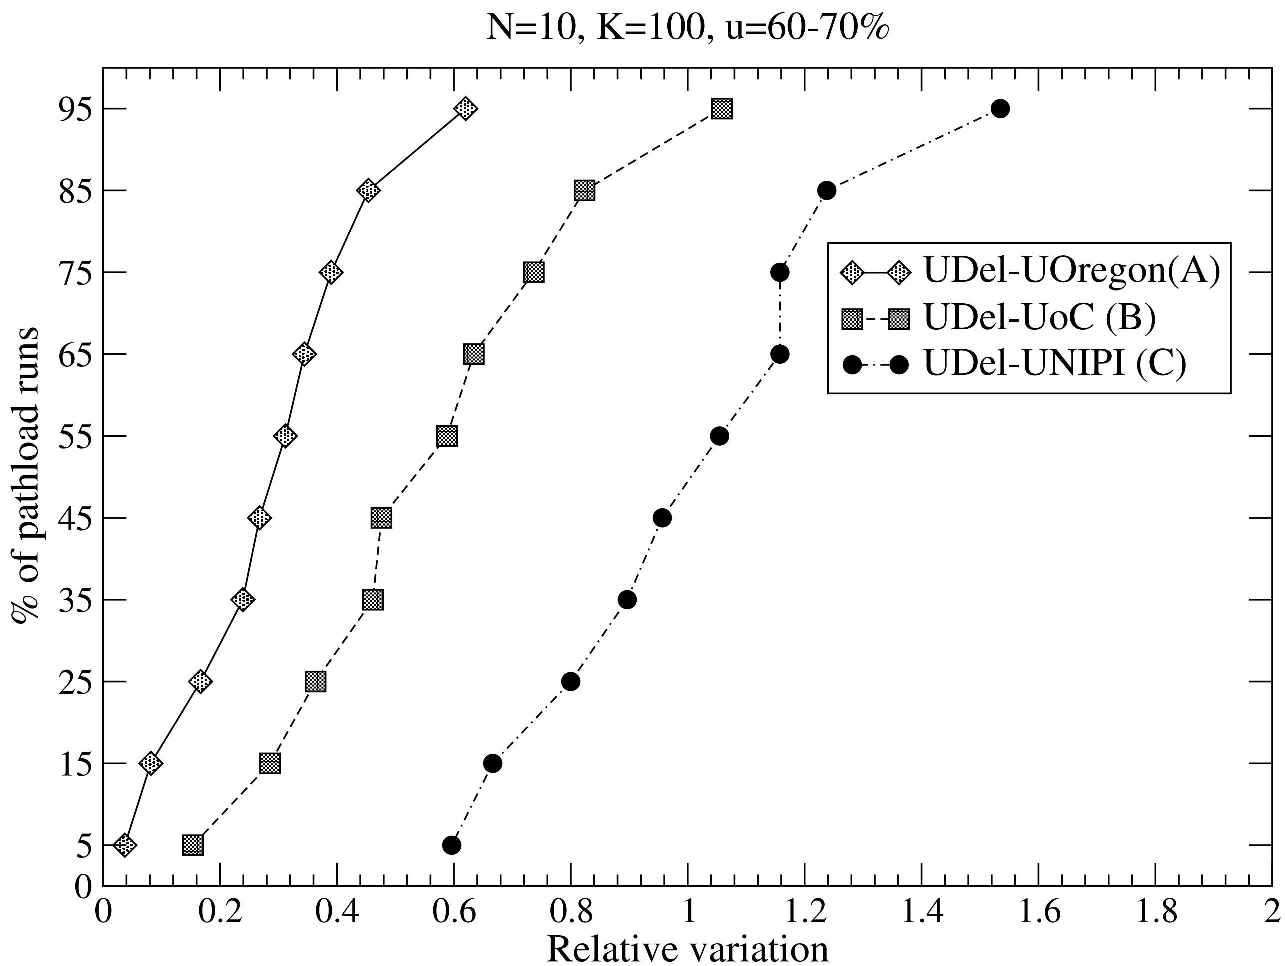

- Traffic aggregation reduces avail-bw variability

Variability and stream duration

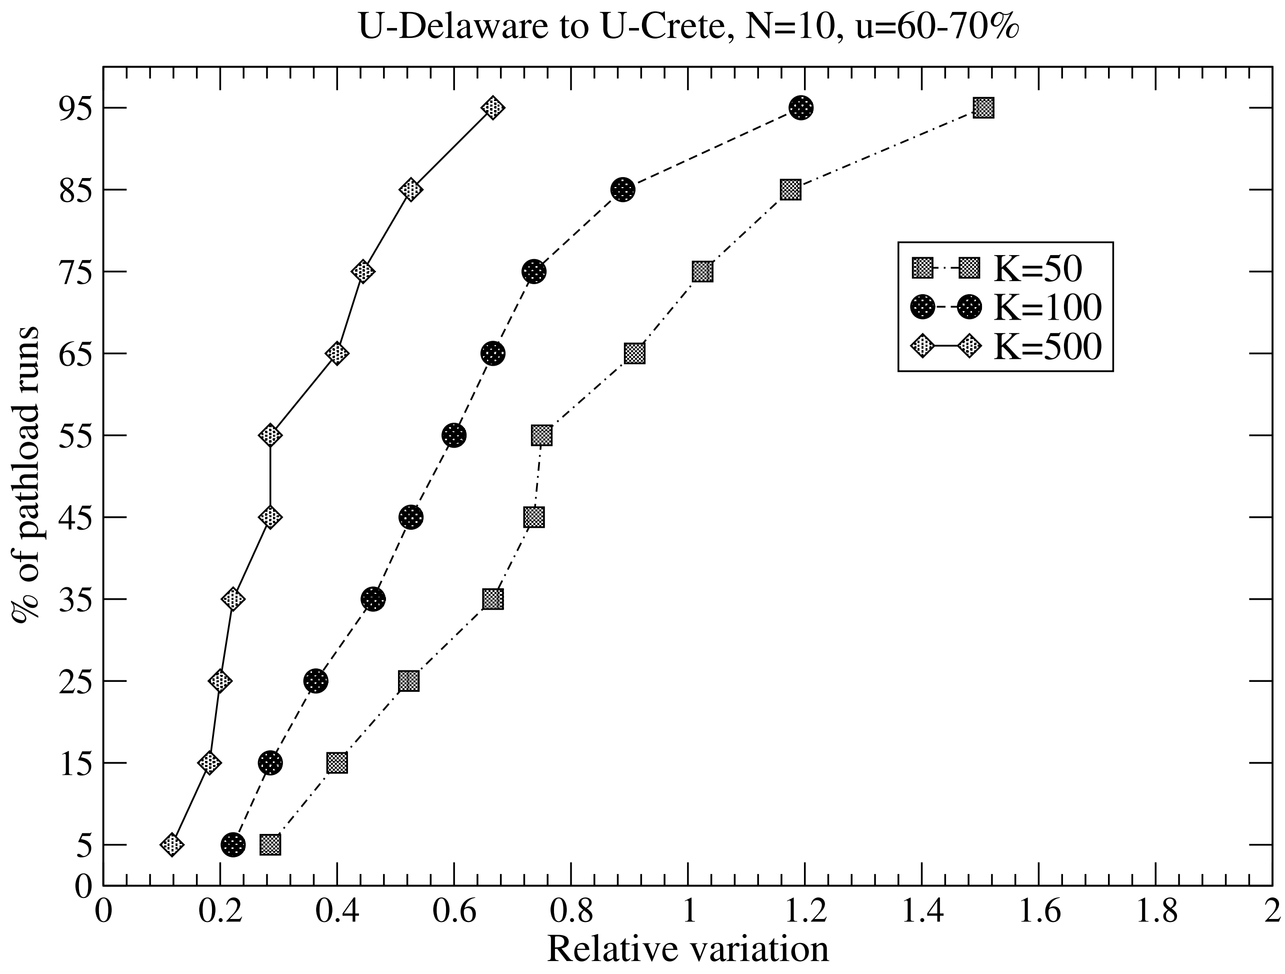

- Variability decreases as stream duration (averaging timescale) increases

Summary

- Avail-bw estimation has numerous applications
- SLoPS: fast, accurate, and non-intrusive measurements
- Implemented in *pathload* (to be released in Spring'02)
- Evaluation of avail-bw variability using *pathload*

- Future work: incorporate avail-bw estimation in transport, QoS, and routing

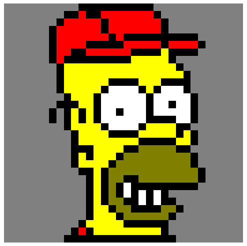

Thank you!