Evaluating pathrate and pathload with realistic cross-traffic

Ravi Prasad
Manish Jain
Constantinos Dovrolis
(ravi, jain, dovrolis@cc.gatech.edu)
College of Computing
Georgia Institute of Technology

Background

- Pathrate
  - Estimates path capacity
  - Based on packet pair/train dispersion
    - Packet pair estimates: Set of possible capacity modes
    - Packet train estimates: ADR=Lower bound on capacity
      - Capacity = (Strongest and narrowest mode > ADR)
- Pathload
  - Estimates path available bandwidth (avail-bw)
  - Based on one-way delay trend of periodic streams
  - Reports a range of avail-bw
    - Corresponds to variation, measured in stream duration
- http://www.pathrate.org

Motivation

- Recent studies pointed towards poor accuracy of these tools
  - http://www.caida.org/outreach/presentations/2003/bwest0308/doereview.pdf
  - *A measurement study of available bandwidth estimation tools*. Strauss et. al. IMC 2003
- Our objective: re-evaluate accuracy of both tools
  - Wide range of cross-traffic load
  - Realistic cross-traffic
  - Completely monitored testbed (no guessing!)

Outline

- Describe test methodology
  - Testbed
  - Cross-traffic type
- Show accuracy results
  - 100Mbps path
  - 1Gbps path
  - With Iperf cross-traffic
- Explaining inaccuracies with Iperf cross-traffic
- Conclusions

Testing methodology

- Used local testbed
  - Complete knowledge of path properties
    - Capacity
    - Available bandwidth
  - Complete control of cross-traffic
    - Rate
    - Type (TCP vs UDP vs trace-driven)

Testbed

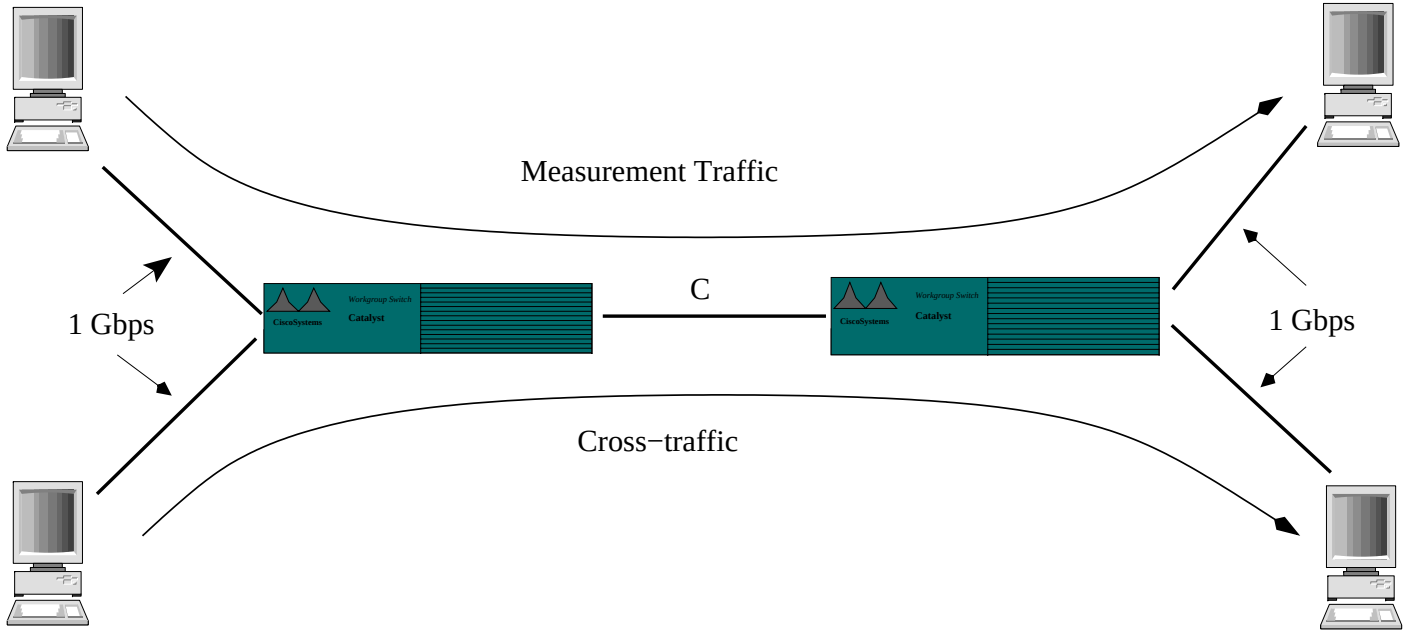

- Narrow link capacity *C* = 100Mbps or 1Gbps

Cross traffic

- Trace-driven cross-traffic generation:
  - NLANR traces
    - OC-3, OC-12, OC-48
      - Trace information at the end of the talk
  - Packet size distribution
    - Unmodified
  - Packet interarrivals
    - Either, scaled to achieve desired cross-traffic throughput
    - Or, unmodified
- Iperf-based cross-traffic
  - Single TCP stream
  - UDP stream

Results

FastEthernet: Traces with scaled interarrivals

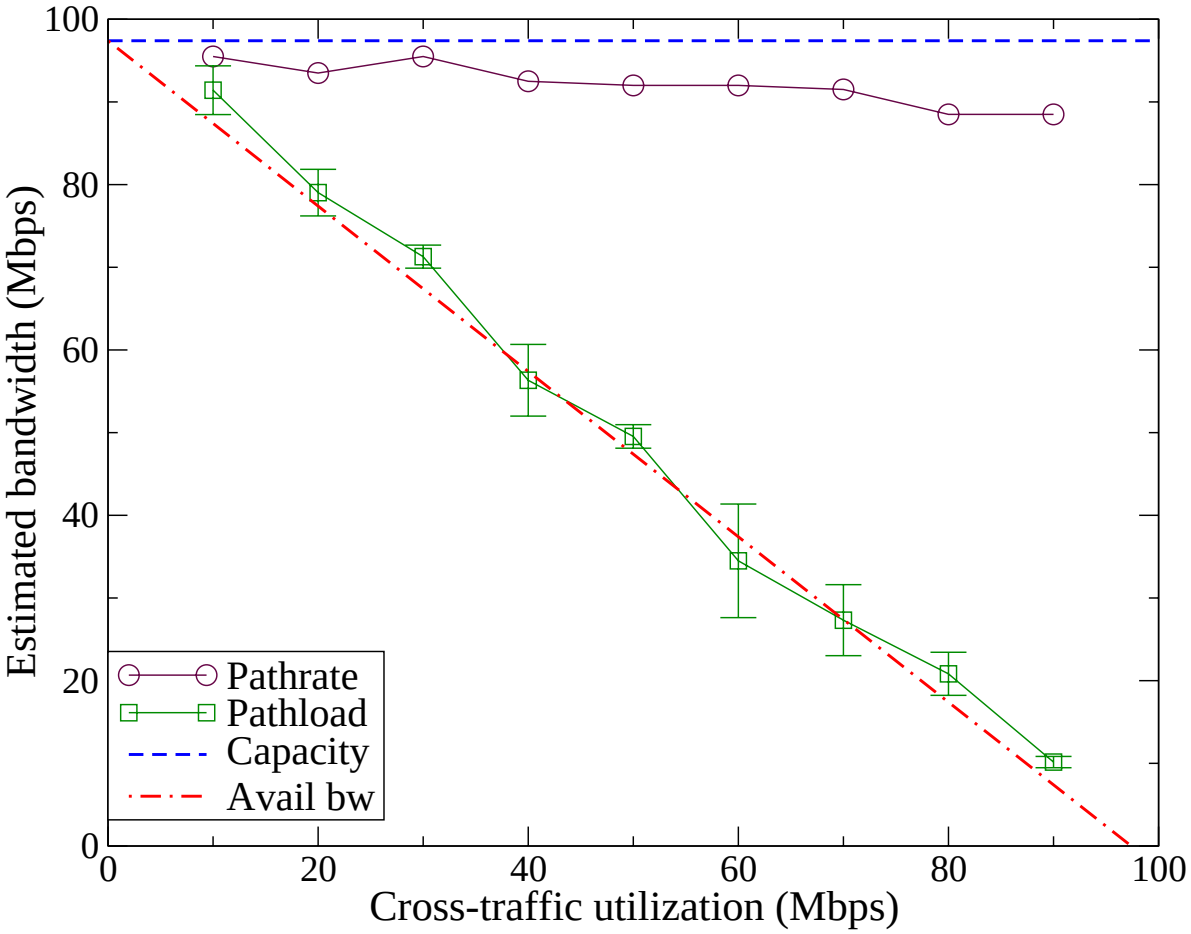

FastEthernet: Traces with unmodified interarrivals

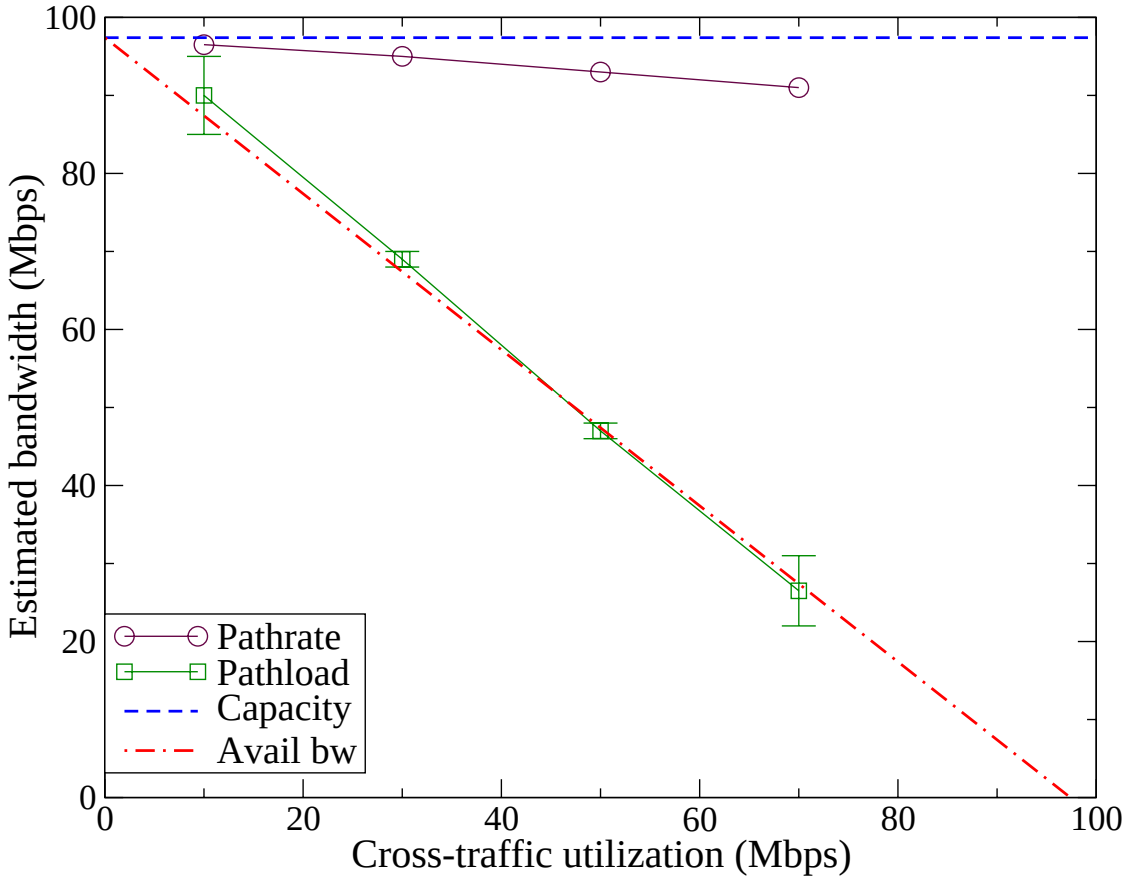

Gigabit path: Traces with scaled interarrivals

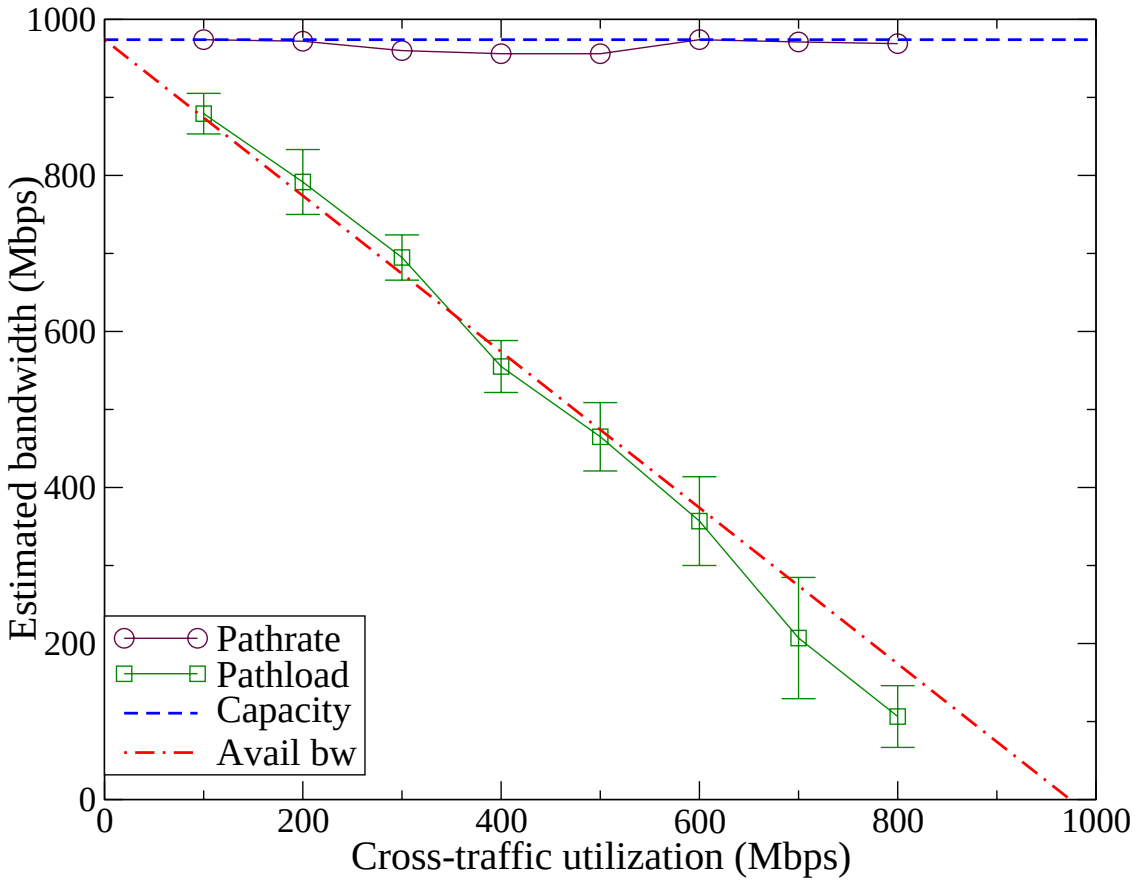

Unrealistic cross-traffic

- Single stream TCP
  - Entire window appears as burst at beginning of RTT
  - Minimum averaging interval: RTT

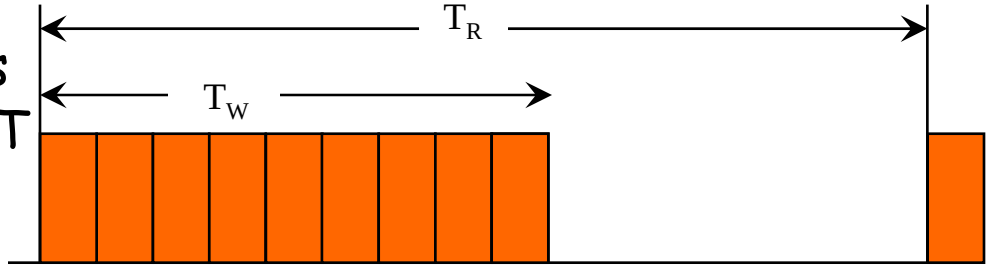

- UDP periodic stream
  - Packet size: L
  - Rate: R
  - Dispersion: L/R
  - Utilization ρ = R/C

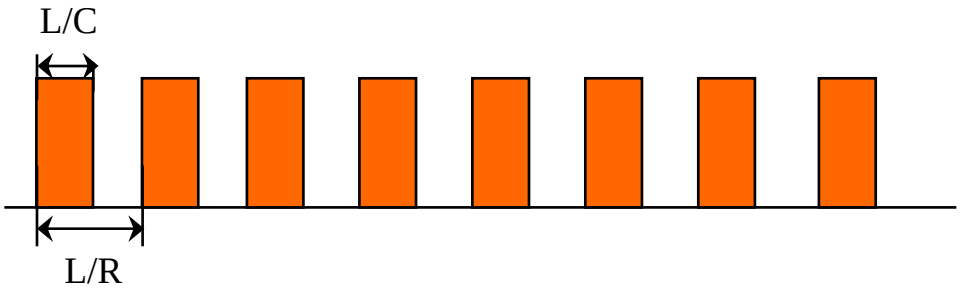

Gigabit Path: Iperf Periodic UDP

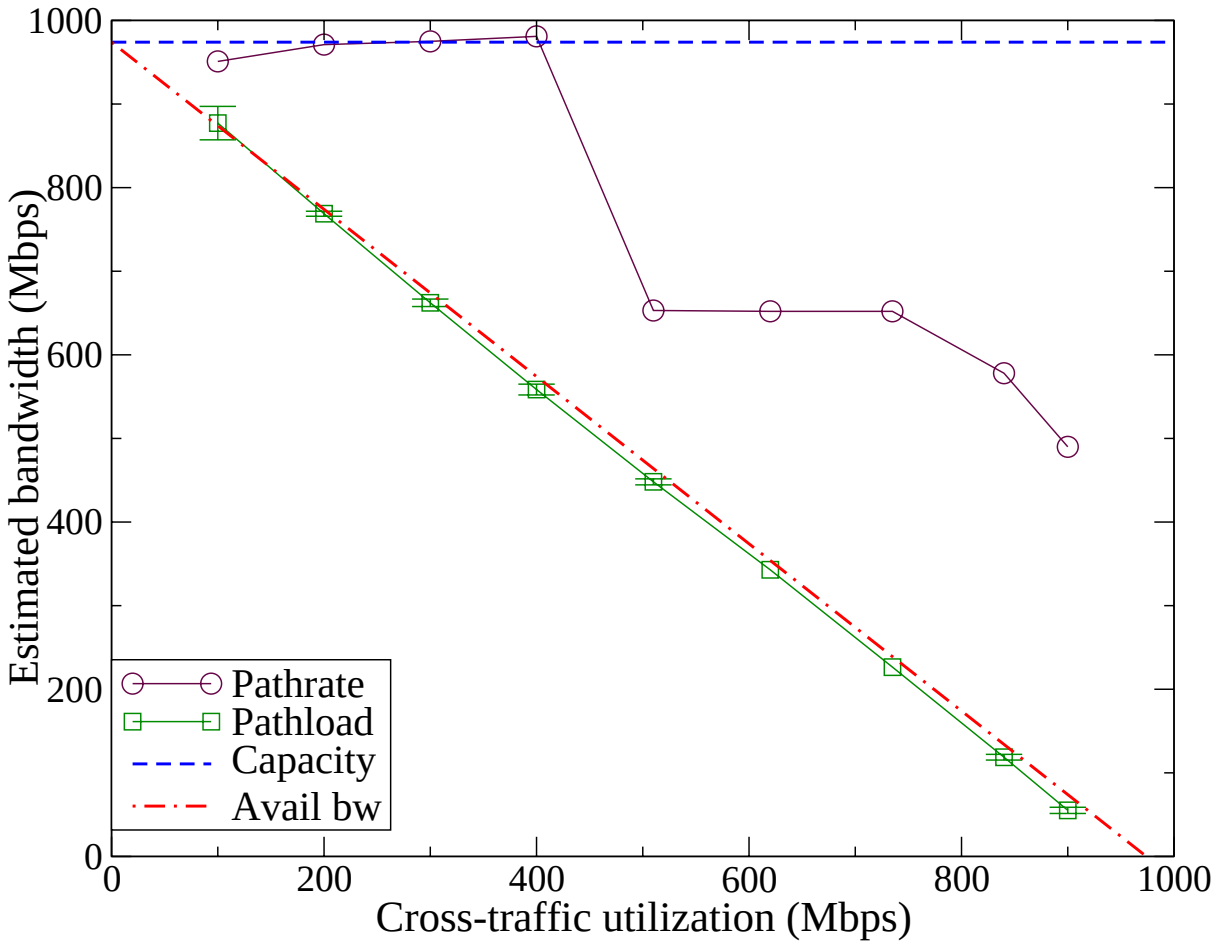

Pathrate under unrealistic traffic

- Seeks some "off" time periods of duration larger than L/C
  - L: Probe size
- TCP traffic
  - Off period $T_O = T_R - T_W - L/C$
    - Correct capacity estimate when $T_O > L/C$
- UDP periodic traffic
  - If $\rho < 0.5$ then $T_O > L/C$
  - Else, underestimation

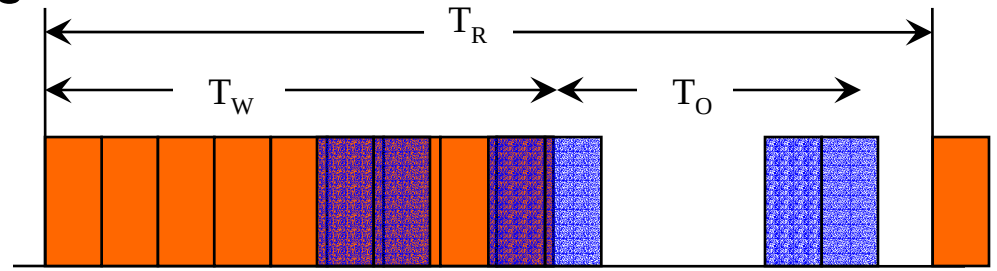

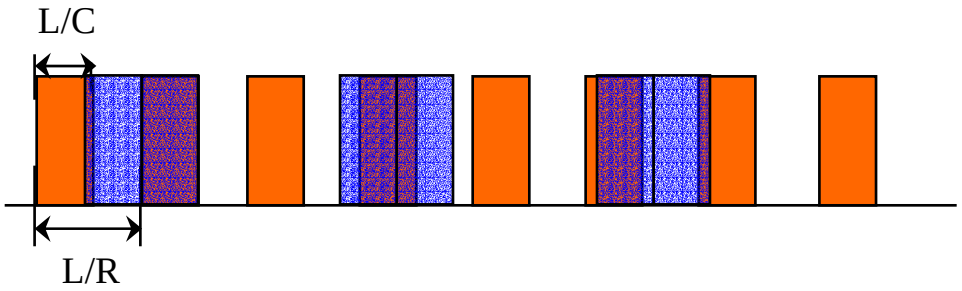

Gigabit path: Iperf single stream TCP

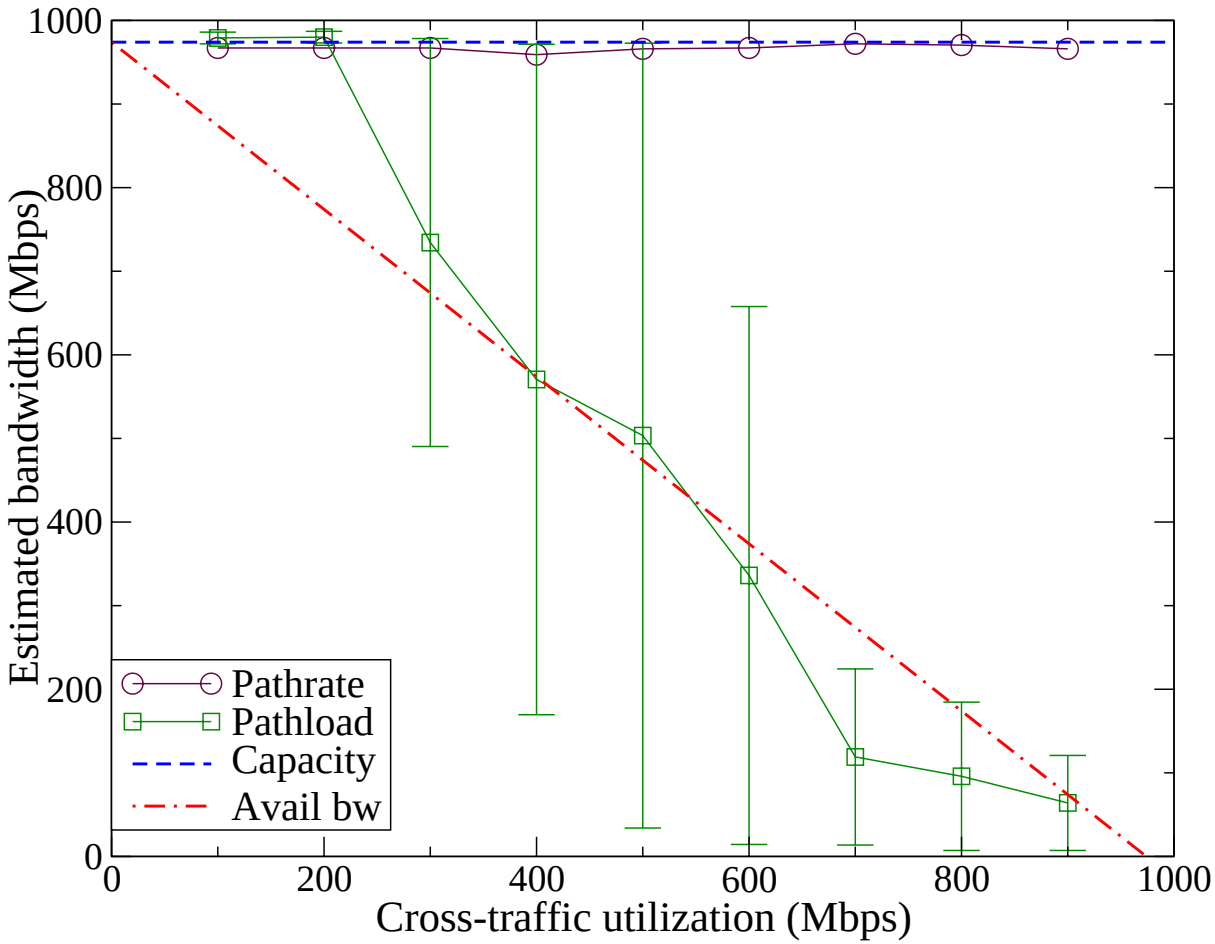

Pathload under unrealistic traffic

- Samples avail-bw in stream duration (T_S)
- TCP traffic
  - Avail-bw averaging period T_R
  - $T_S << T_R$ results in wide Avail-bw range estimate

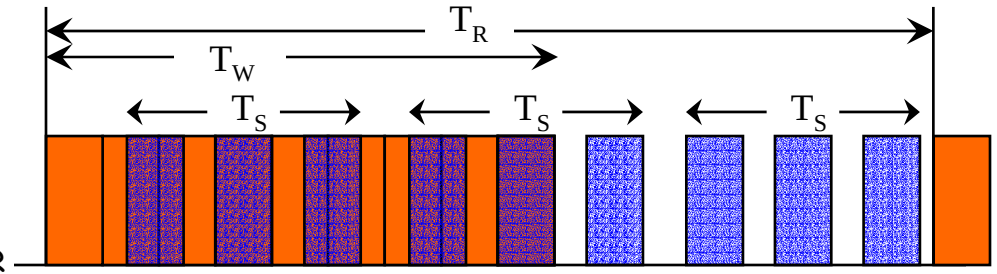

- UDP periodic traffic
  - Avail-bw averaging period L/R
  - T_S = 100 x L/C > L/R
  - Correct avail-bw range estimate

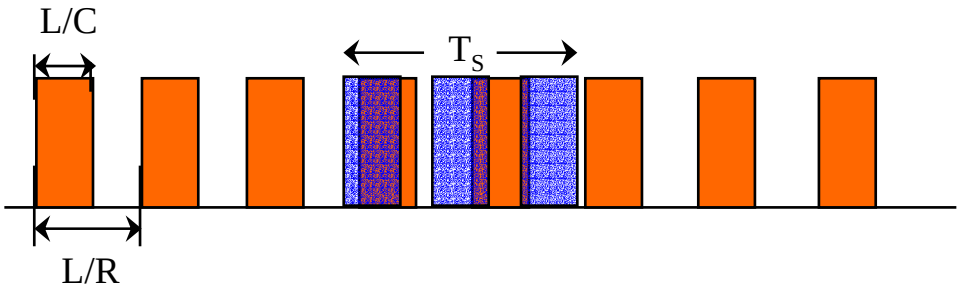

Conclusions

- Type of cross-traffic is important for bandwidth estimation tools
- Pathrate and pathload perform well with realistic cross-traffic
- Simulated traffic does not capture:
  - Packet size distribution
  - Interarrival distribution
  - Correlation structure

Trace identifiers

- OC3 : MEM-1070464136-1,
  COS-1070488076-1,
  BWY-1063315231-1,
  COS-1049166362-1
- OC12: MRA-1060885637-1
- OC48: IPLS-CLEV-20020814-093000-1
- We greatly appreciate the availability of traces from NLANR PMA project. The NLANR PMA project is supported by National Science Foundation Cooperative agreement nos. ANI-0129677 (2002) and ANI-9807479 (1998).