A test-bed investigation of QoS mechanisms for supporting SLAs in IPv6

Vasilios A. Siris and Georgios Fotiadis

University of Crete and FORTH
Heraklion, Crete, Greece

vsiris@ics.forth.gr

Funded in part by GRNET, through subcontract in 6NET (IST-2001-32603)

Objectives and motivation

- Investigate operation and interaction of QoS mechanisms in IPv6 networks
  - Policing, shaping, queuing
- These mechanisms will be used in customer-provider interface
- Tuning is important for supporting Service Level Agreements (SLAs)
  - Tuning each in isolation not sufficient
- Do they have different performance in IPv6 compared to IPv4?

User-network interface

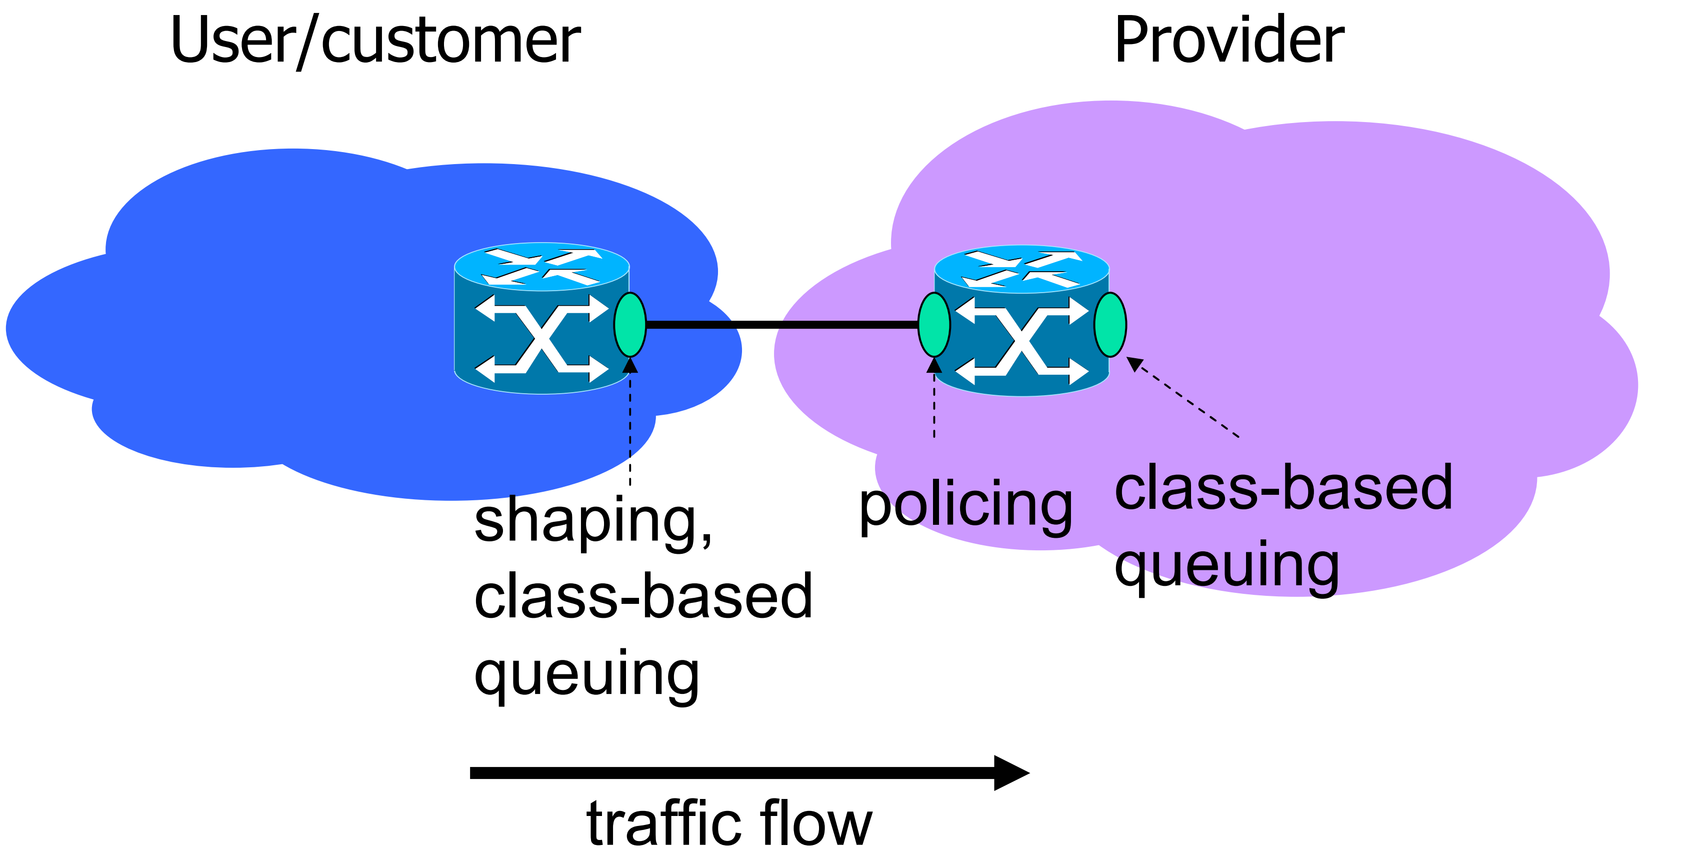

QoS mechanisms investigated

- Cisco Traffic Policing
  - *Committed Access Rate – CAR not supported in IPv6*
- Cisco Generic Traffic Shaping - GTS
- Linux Traffic Policing
- Linux Traffic Shaping
  - Token Bucket Filter – TBF
- Linux Class Based Queuing - CBQ

Token bucket algorithm

- True token bucket (r,b)
- Shaping includes a buffer

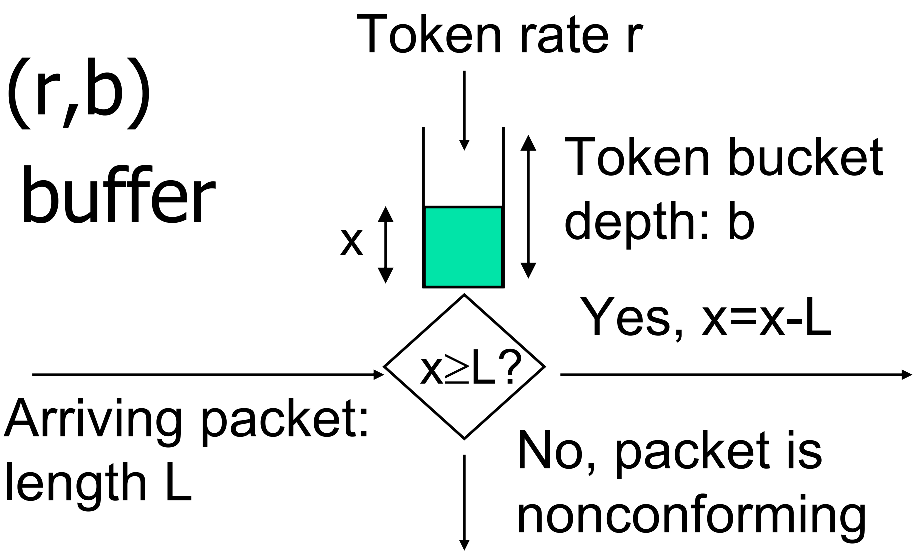

- Implementation differs
  - Mean rate or Committed Information Rate (CIR)
  - Committed Burst Size (Bc)
  - Time interval (=Bc/CIR)

IPv6 test-bed

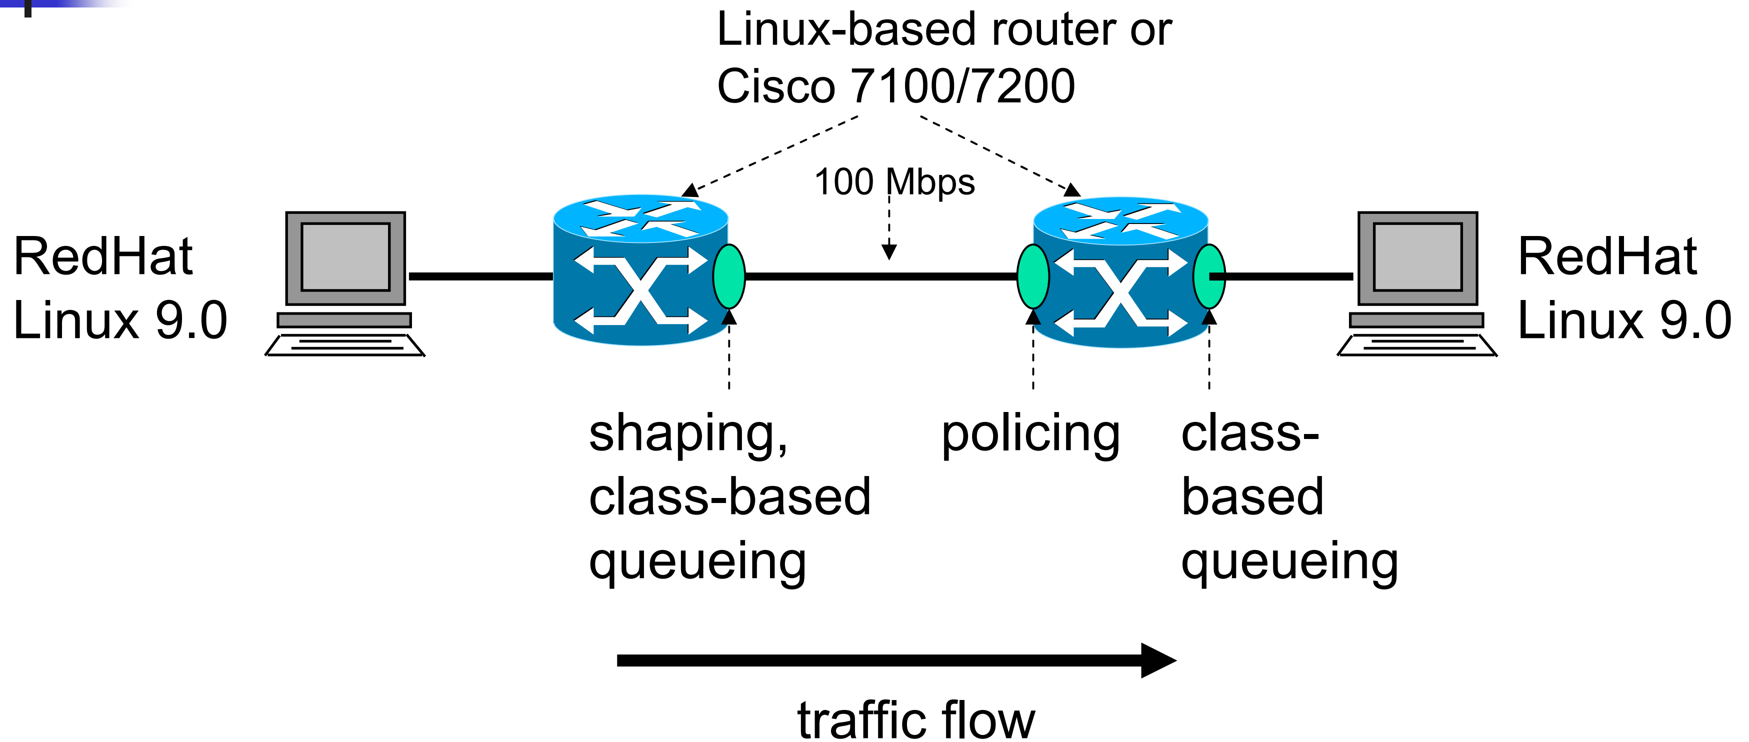

- TCP traffic generated with Iperf 1.7.0 for Linux
- RTT < 4 ms

Policing tests topology

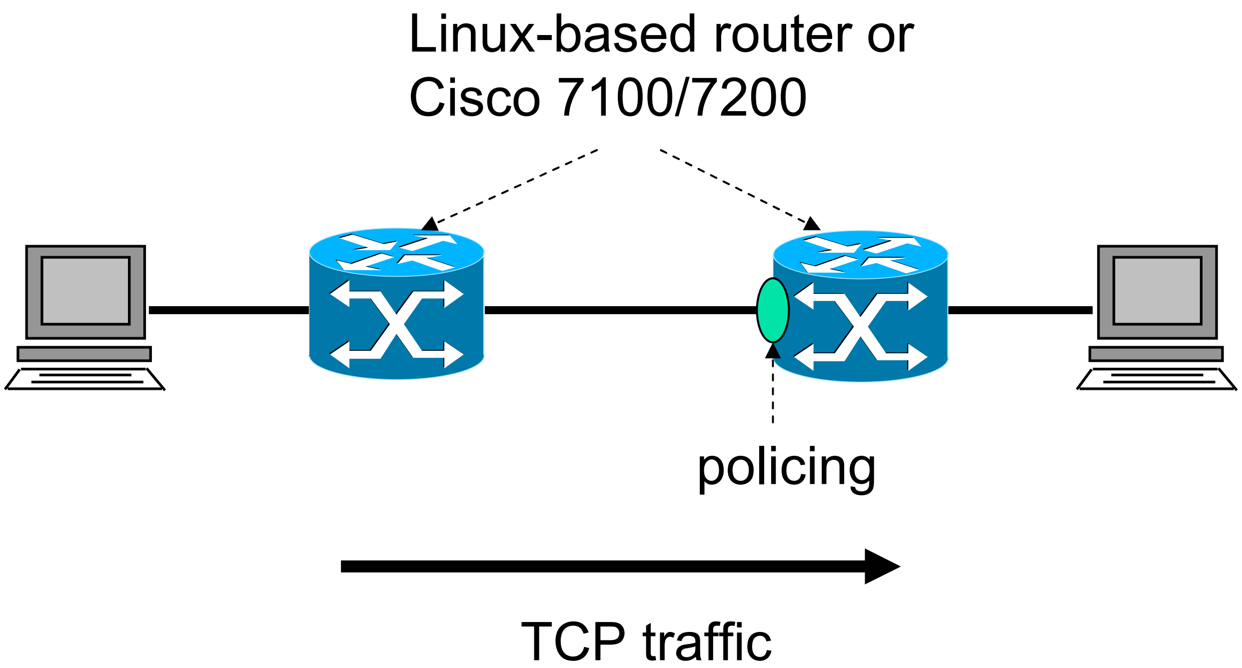

- For fixed policing rate, measure aggregate throughput for different burst sizes

Variable policing rates (1/3)

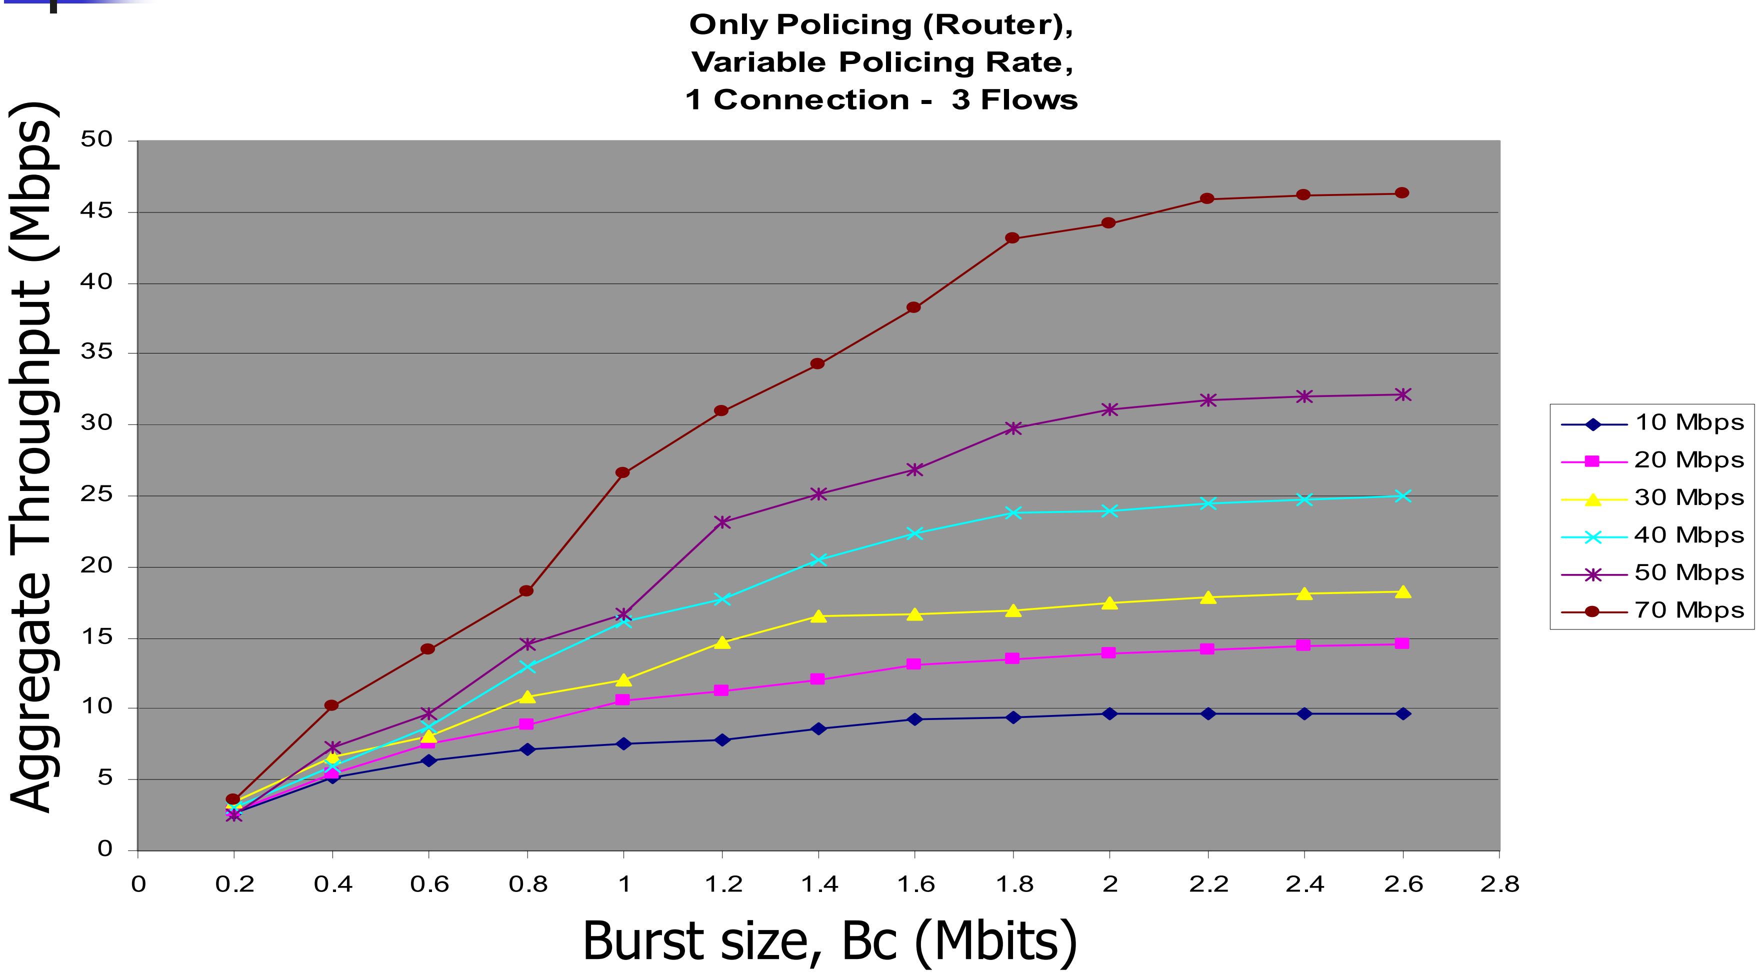

Variable policing rates (2/3)

- Aggregate throughput increases with burst size, but is always < policing rate
- For higher rates throughput is proportionally smaller
  - Rate = 10 Mbps -> Throughput = 9.7
  - Rate = 70 Mbps -> Throughput = 46
- Experiments for both Cisco & Linux routers gave identical results

Variable policing rates (3/3)

- Throughput versus burst size exhibits a "knee" at Bc*, which is approximately:
  - If Rate < 40Mbps,

    Bc* = 0.6 Mbit + 0.03sec*Rate
  - If Rate > 40Mbps,

    Bc* = 1.5 Mbit + 0.01sec*Rate
- Burst Size > RTT*Rate not sufficient!

Variable number of flows (1/2)

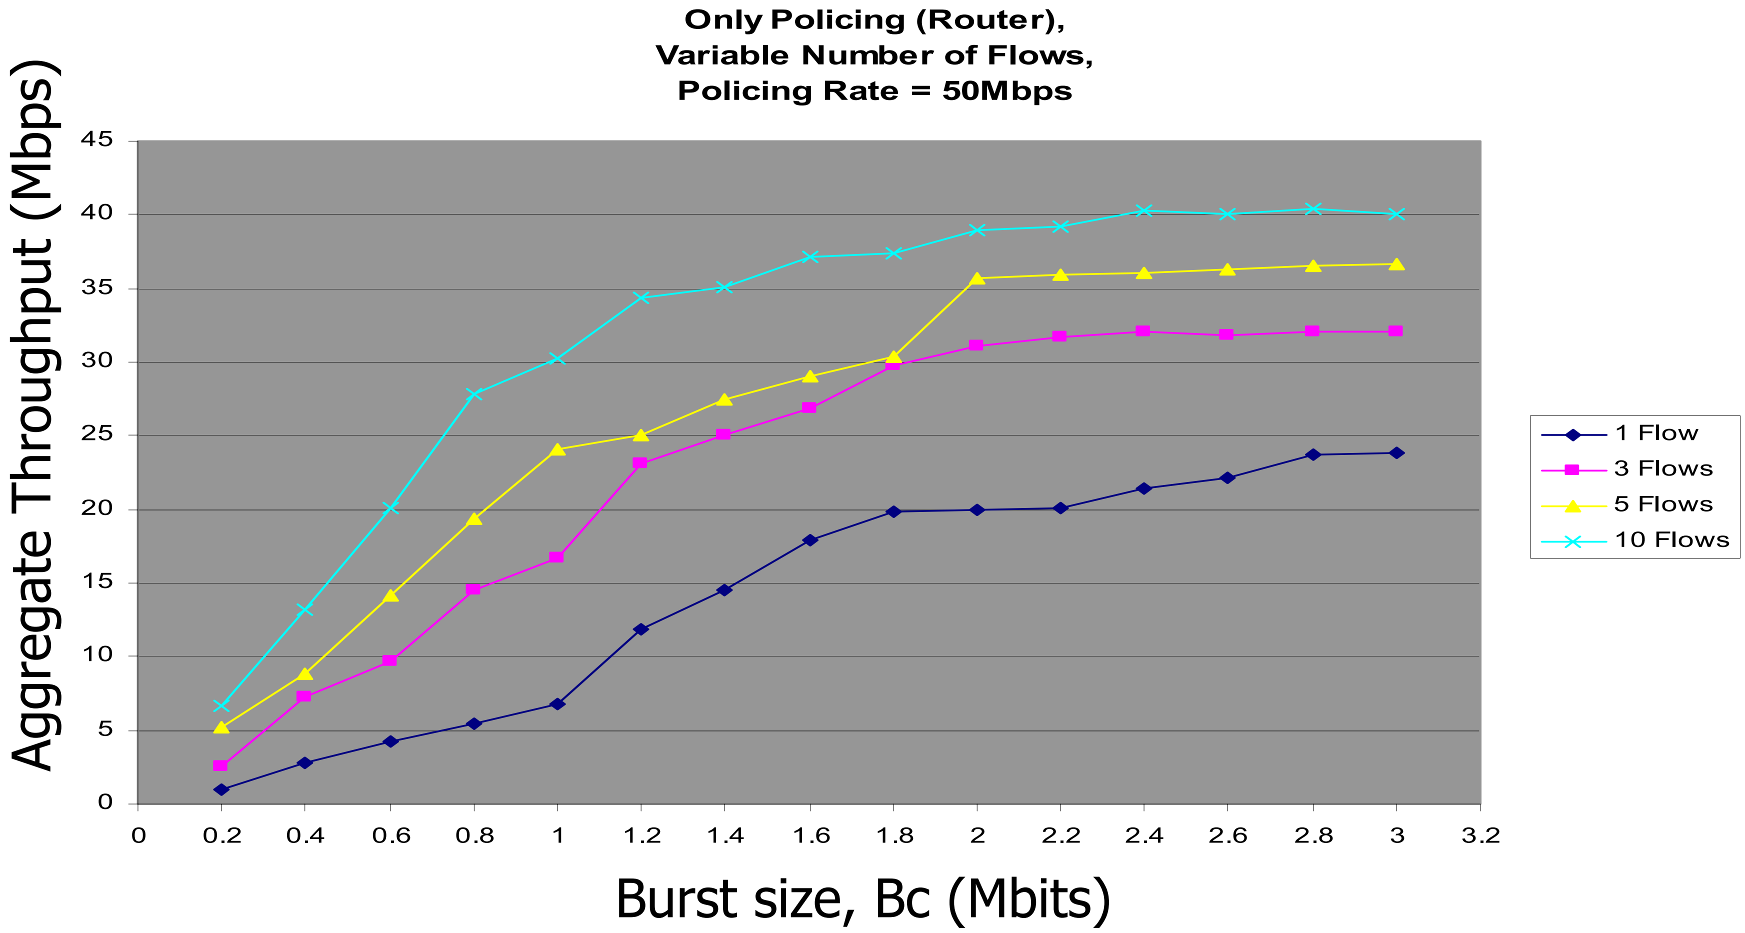

Variable number of flows (2/2)

- Optimum burst size (“knee” value) is independent of the number of flows
- Aggregate throughput increases with number of flows (multiplexing)
- Similar results for Cisco & Linux routers

Interaction of policing & shaping

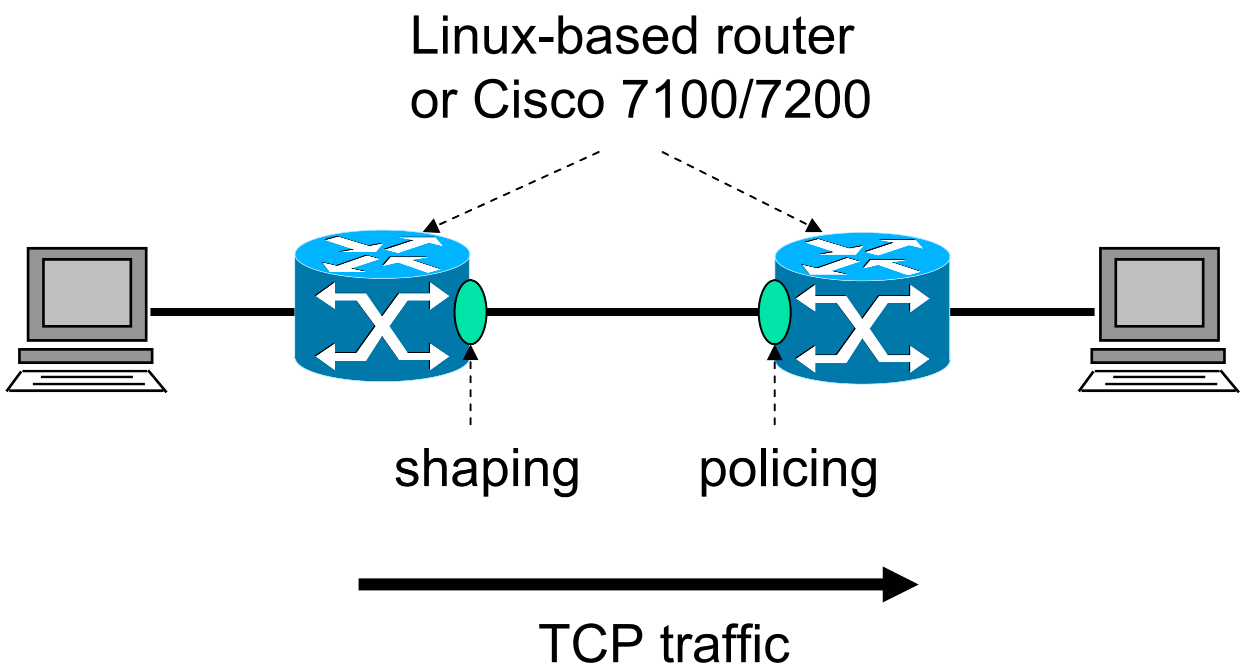

- For fixed rate & burst size of policer and rate of shaper, measure aggregate throughput for different burst sizes of shaper

Throughput vs. burst size (1/4)

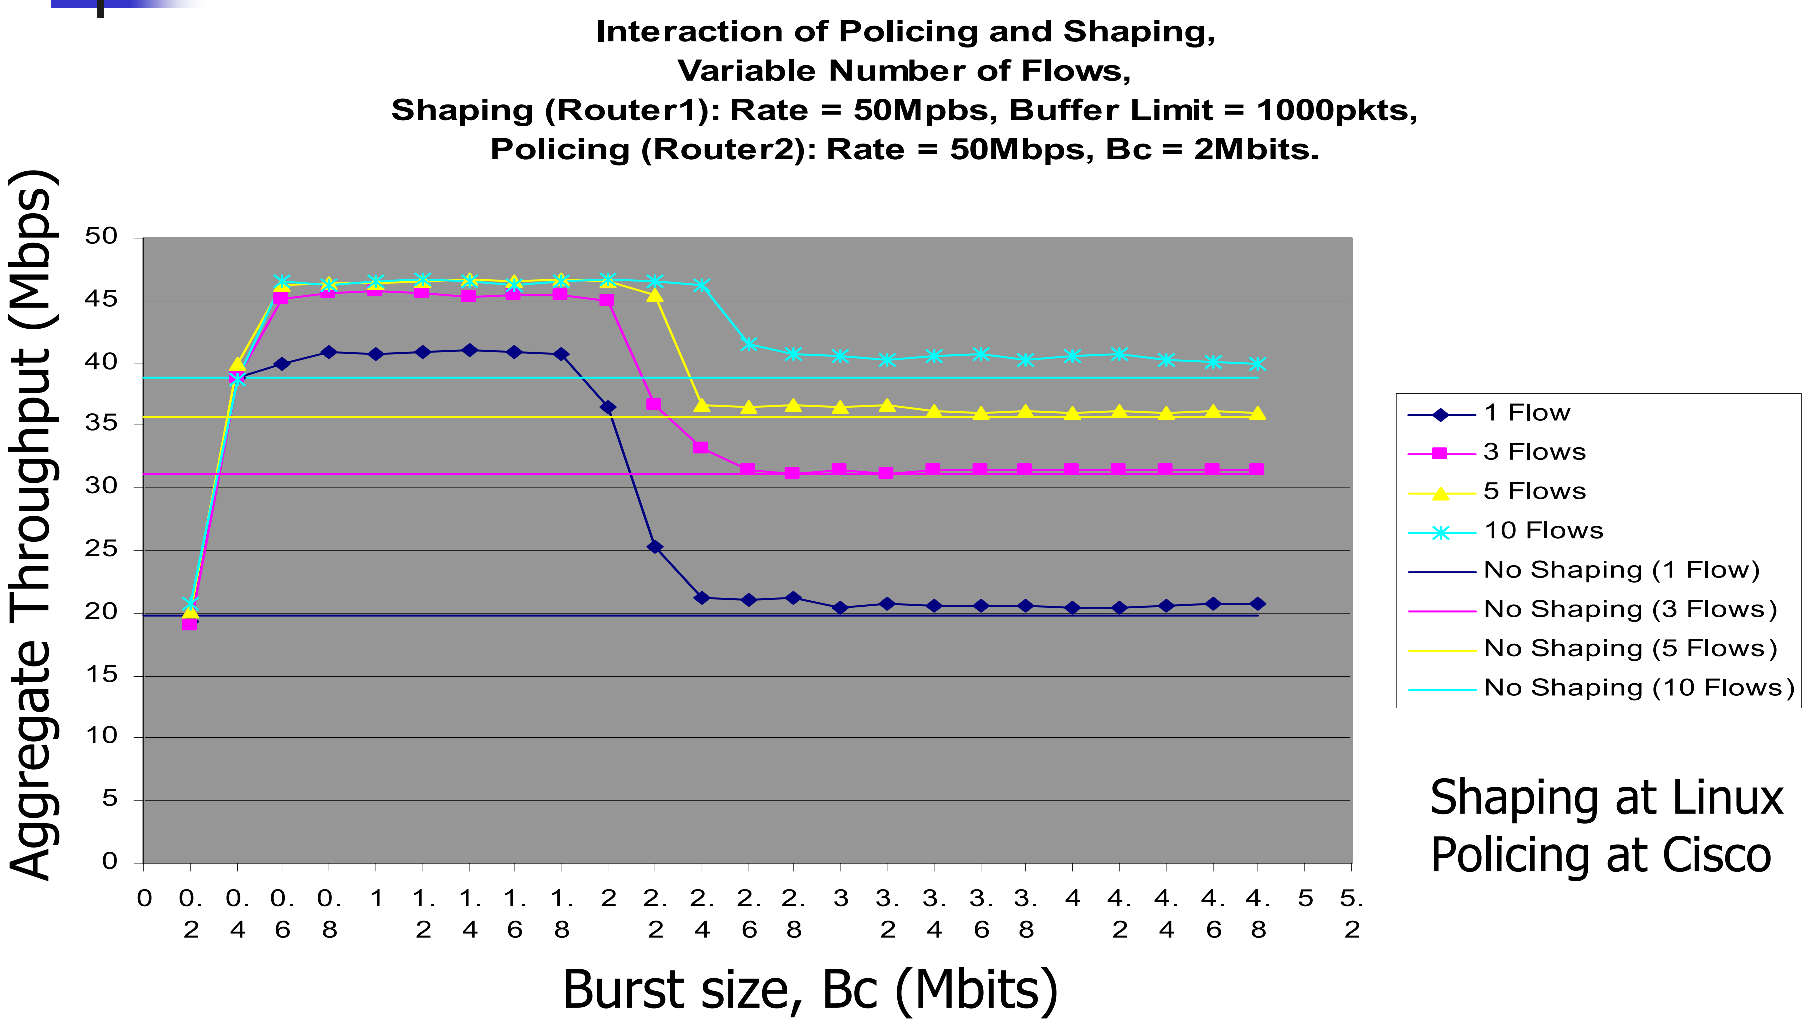

Shaping at Linux
Policing at Cisco

Throughput vs. burst size (2/4)

- Appropriate shaper burst size can lead to significant throughput increase
  - Throughput is always higher with shaping
  - Setting burst size of shaper = burst size of policer not always optimal
- Low throughput for very small burst sizes
  - Due to overflow at shaper's queue
- But also, very large burst size gives low throughput
  - Very large burst sizes allow large bursts which lead to multiple drops at the policer

Throughput vs. burst size (3/4)

- For larger # of flows there is larger range of burst sizes that achieve maximum throughput

Throughput vs. burst size (4/4)

No qualitative difference when using two Linux or Linux & Cisco routers

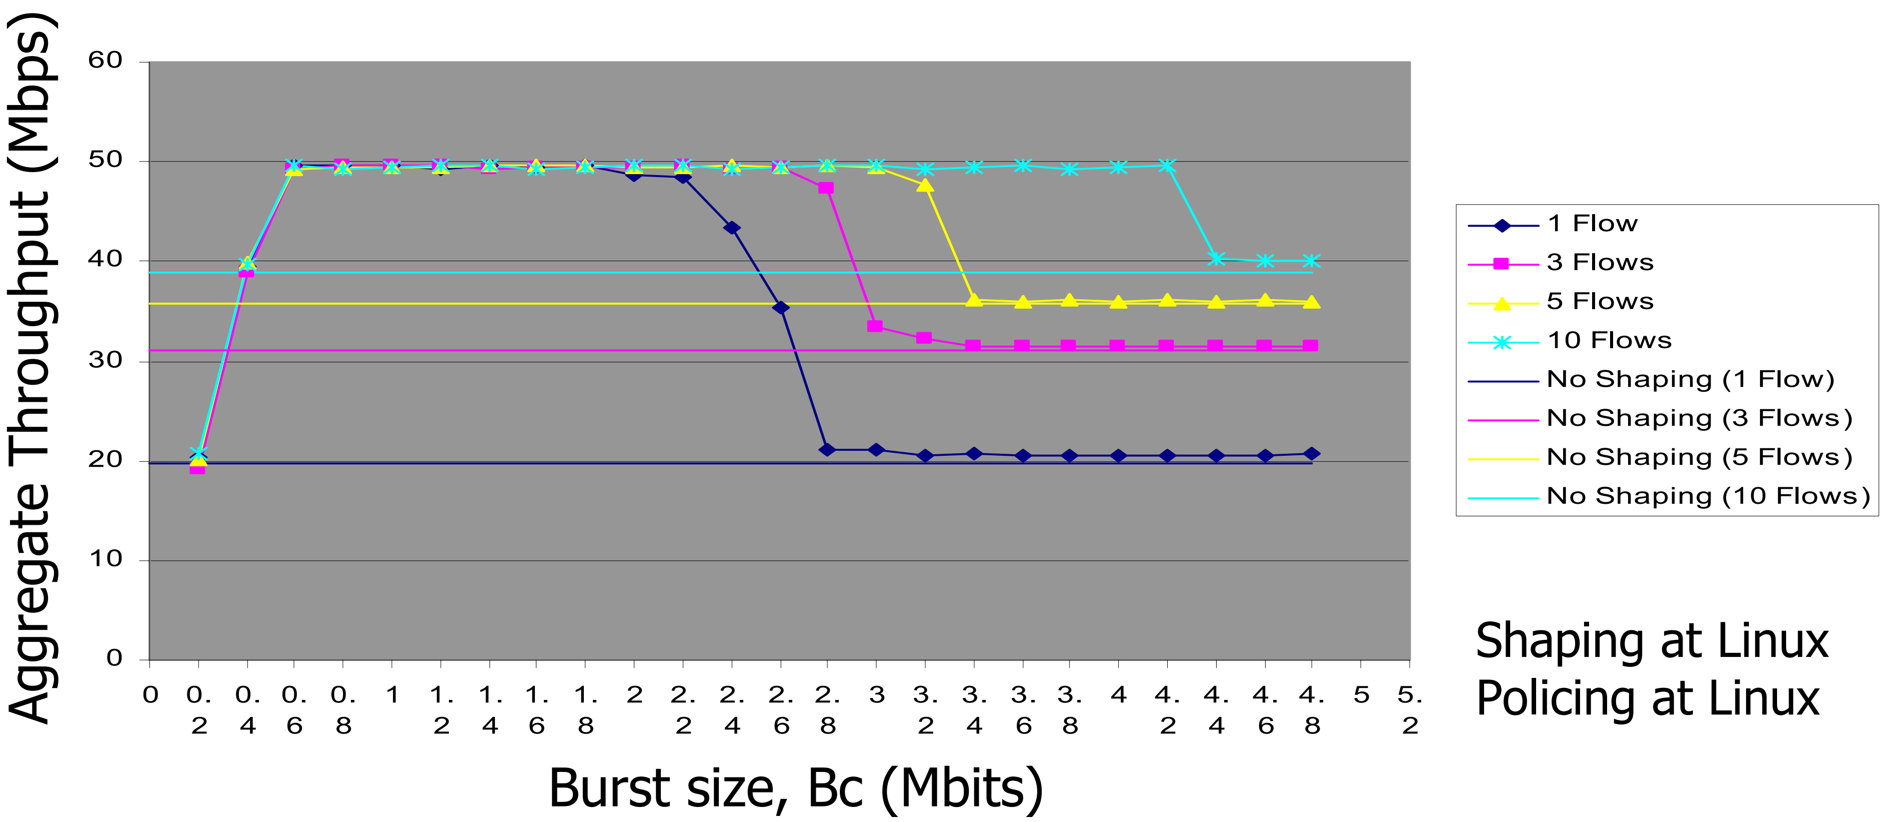

Throughput vs. shaping rate (1/2)

Shaping at Linux
Policing at Cisco

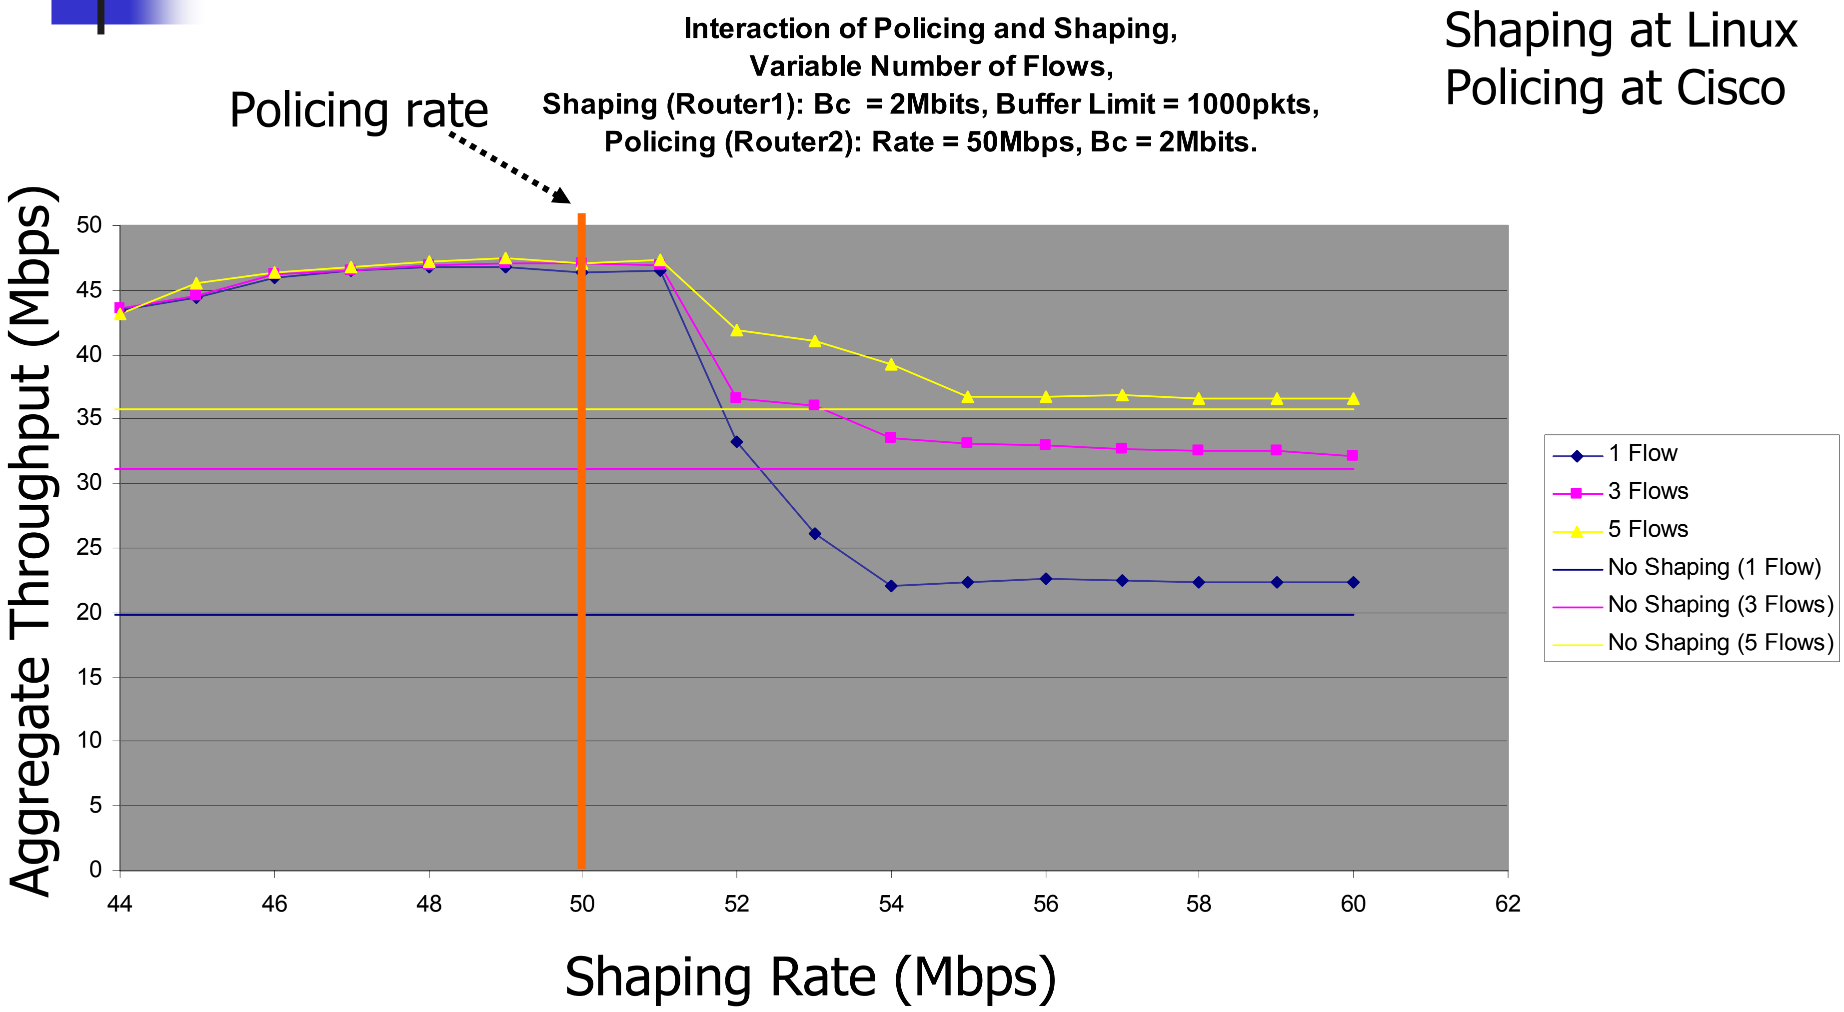

Throughput vs. shaping rate (2/2)

- Throughput maximum for shaping rate a little higher than policing rate
- Very high shaping rates leads to throughput degradation, due to bursts that are allowed by the shaper

Delay jitter tests (1/3)

- Policing/Shaping tests showed there is a range of shaper burst sizes for which throughput is maximized
- Which burst size is best?

- Answer: consider delay

Delay jitter tests (2/3)

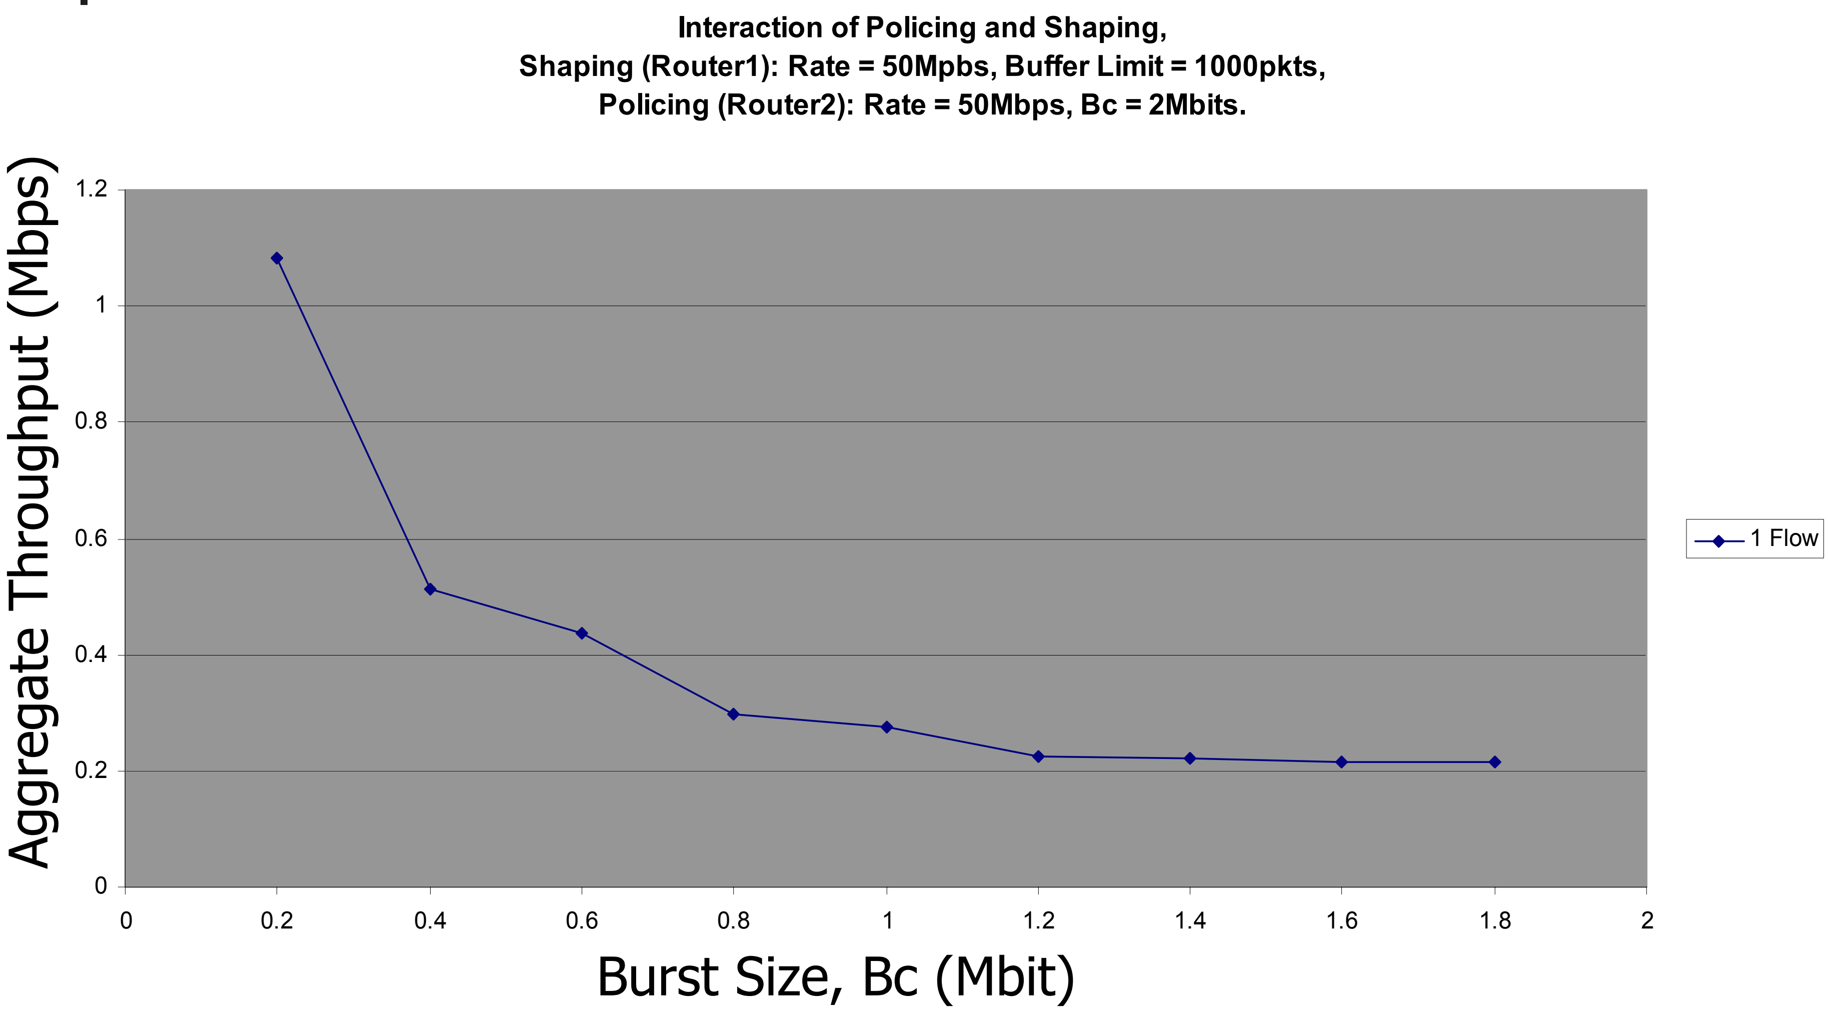

Delay jitter tests (3/3)

- Jitter decreases when burst size increases
- Optimal burst size is one for which delay jitter is smallest, but is in the range of values that maximize throughput

Interaction of Policing and Class Base Queuing

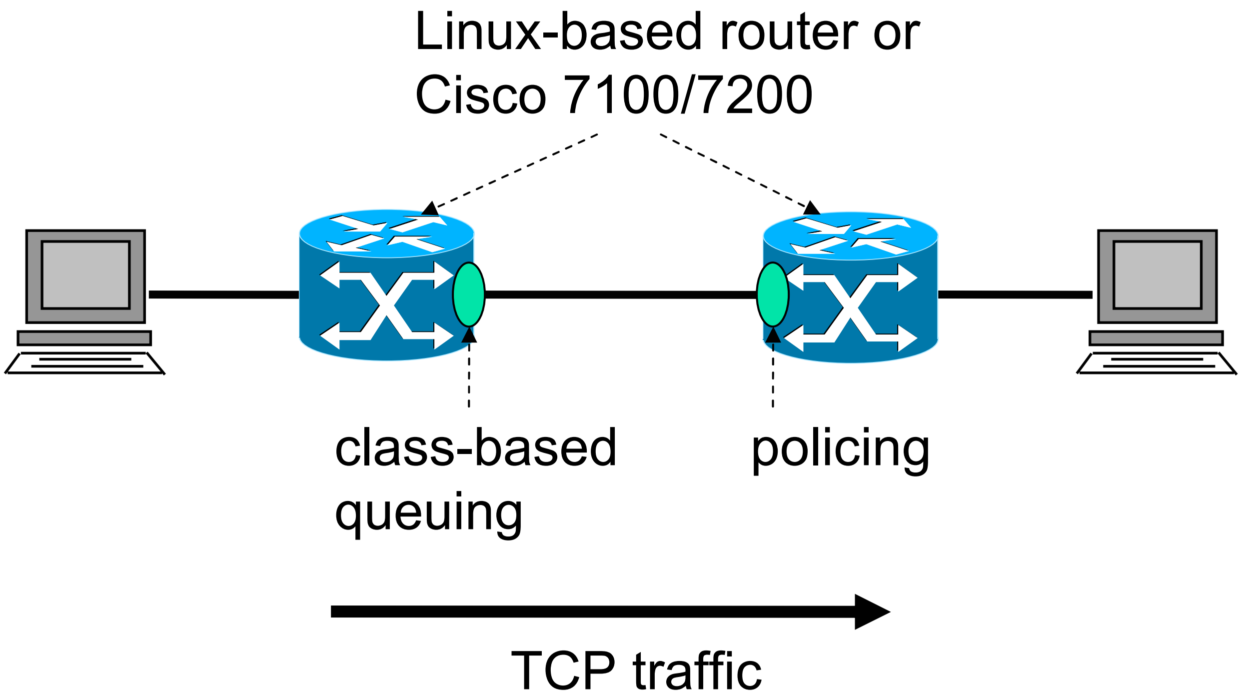

- For fixed rate & burst size of policer and rate of CBQ, measure aggregate throughput for different maxburst values

Variable *maxburst* at CBQ (1/2)

CBQ at Linux
Policing at Cisco

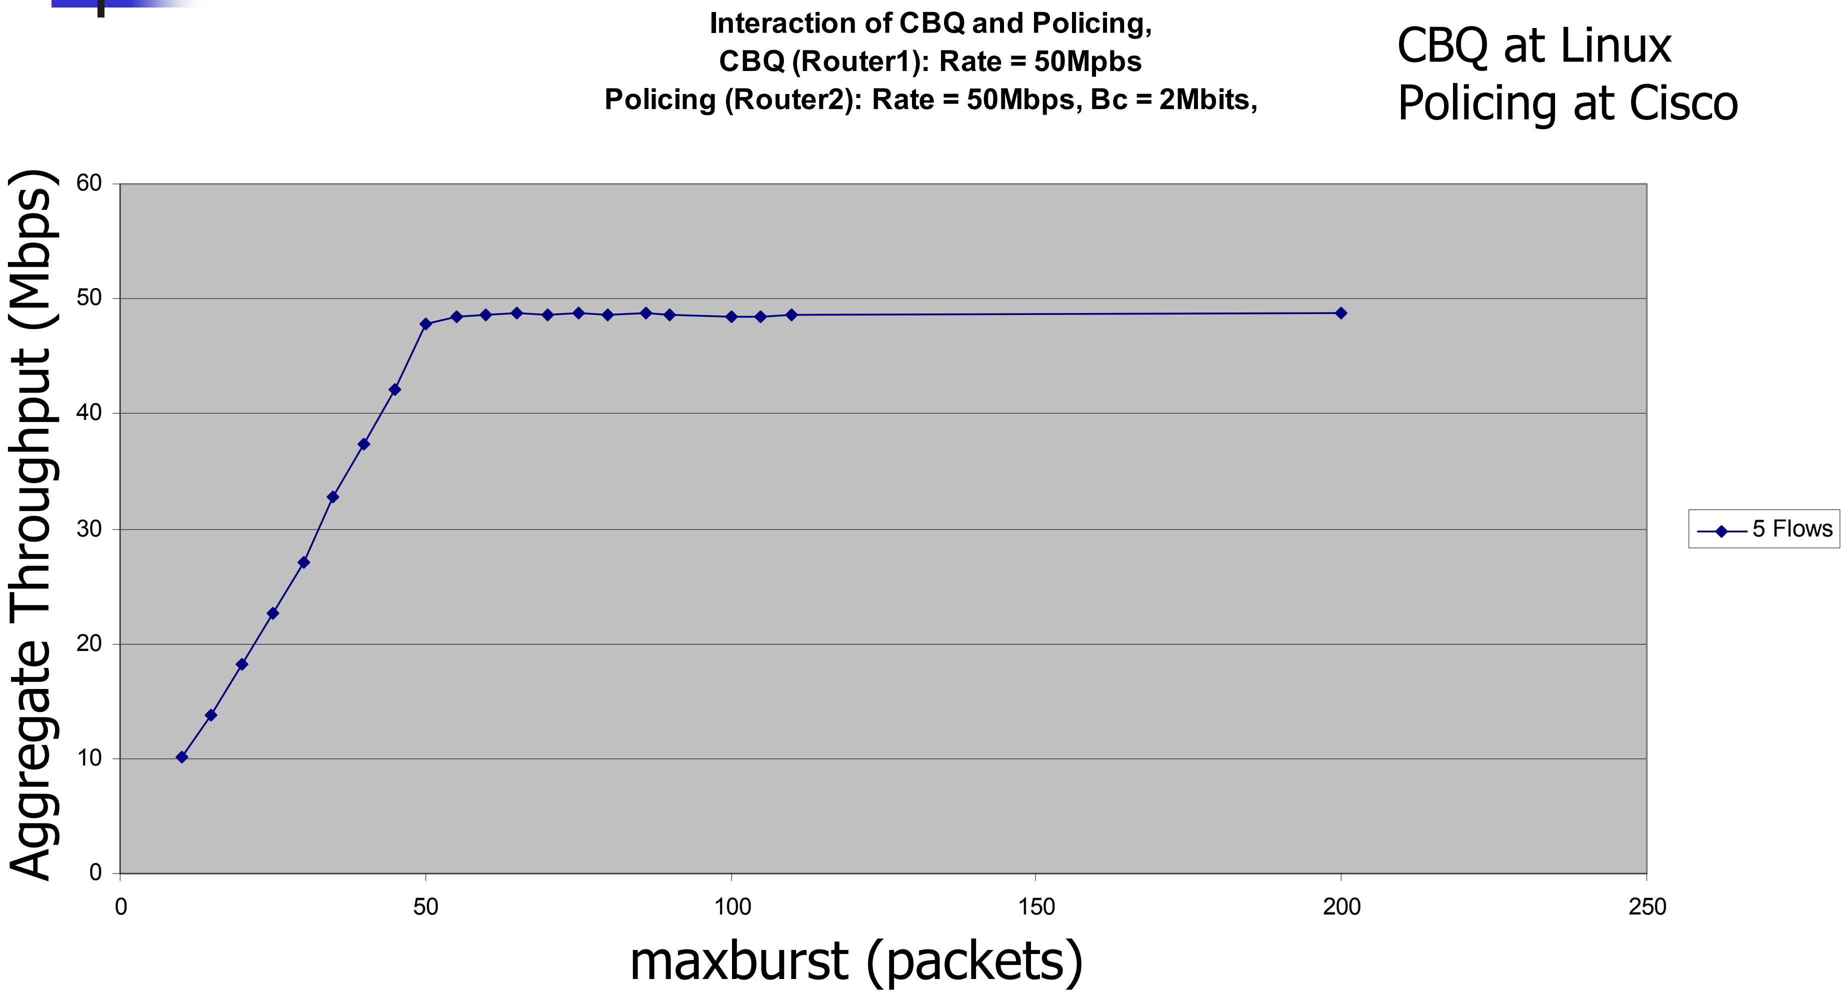

Variable *maxburst* at CBQ (2/2)

- Maximum throughput achieved with appropriate *maxburst* value
- Low throughput for small *maxburst*
- Large *maxburst* values do not result in throughput degradation
  - maxburst has different effect than burst size in shaping
  - CBQ performs traffic smoothing

Conclusions and outlook

- Interaction of QoS mechanisms is important for effectively supporting SLAs
- In the presence of policing, shaping can increase aggregate throughput
  - Must appropriately set shaping parameters
- Tuning of parameters should consider throughput & delay
- Similar behavior with Cisco/Linux, IPv6/IPv4
- Further research directions:
  - Experiments in wide area (larger RTT)
  - Experiments in wireless networks

A test-bed investigation of QoS mechanisms for supporting SLAs in IPv6

Vasilios A. Siris and Georgios Fotiadis

University of Crete and FORTH
Heraklion, Crete, Greece

vsiris@ics.forth.gr

Funded in part by GRNET, through subcontract in 6NET (IST-2001-32603)