Interaction of ATM and IP congestion control

Henning Schulzrinne
Dept. of Computer Science
Columbia University
schulzrinne@cs.columbia.edu

(joint work with Dorgham Sisalem, GMD Fokus)

IP+ATM Workshop
St. Louis, MO

November 13-14, 1996

Overview

- TCP next to ABR
- TCP over ABR
- “ABR-friendly” TCP
- audio/video

TCP next to ABR

- "filler" scenario: TCP over UBR, with ABR, VBR, CBR at higher priority
- ➠ bandwidth visible to TCP fluctuates
- losses ➠ likely less TCP throughput

TCP over ABR

ABR:

- $+$: high utilization, fairness
- $+$: low buffering requirements
- $-$: high complexity: forward, backward, EFCI, …
- performance for short or on/off connections?
- not as loss-sensitive as TCP (see Lakshman/Madhow)

Fairness values for generic fairness configuration:

Plain TCP (no ATM)	0.67
TCP over ATM	0.62
TCP over ATM, EPD, RED	0.63
TCP over ATM, EPRCA	0.99

TCP over ABR

TCP loop = ABR loop ➠ backpressure at source ➠

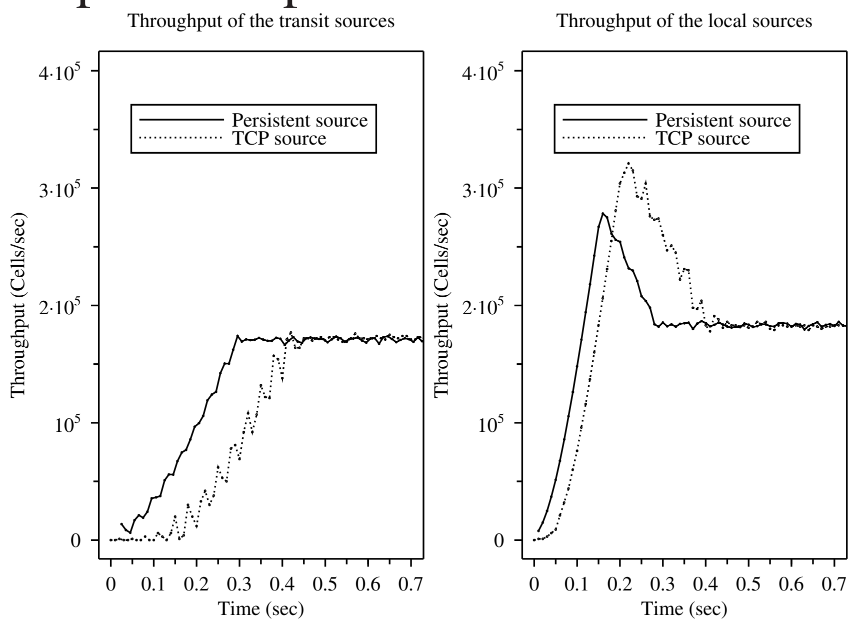

TCP with BCN

- add ATM-like EFCI option:
  - TCP option ➠ layer violation
  - IPv6 header option
- might be less loss-sensitive
- maybe use explicit rate?

TCP with ABR in the middle

200 m access:

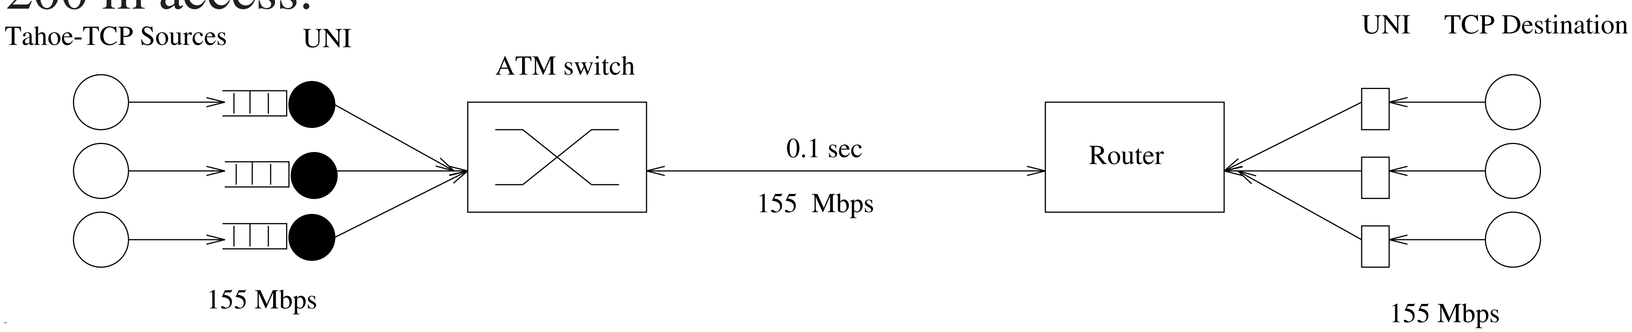

Utilization (%) [without ATM: ≈ 96%]:

buffer	Tahoe	Reno	BCN
5 kB	95	86	99
25 kB	98	90	99

➠ works well for short access links (campus)
Buffering at transition delays bandwidth decrease
Longer access delays ➠ ?

Interoperability of BCN & “classic” TCP

- BCN gets $4\times$ Reno throughput
- ➠ different dropping and marking thresholds
- BCN does not fix distance penalty

Video Traffic

- audio traffic ➠ CBR
- video traffic: CBR + ABR
- ABR ➠ feedback loop $\rightarrow$ change quantization, frame rate (Vickers/Suda, Bolow/Schulzrinne, Ramakrishnan/Lakshman/Mishra)