

97-0607: Simulation Experiments with Guaranteed Frame Rate for TCP/IP Traffic

**Rohit Goyal, Raj Jain, Sonia Fahmy,
Bobby Vandalore, Shivkumar Kalyanaraman**

The Ohio State University

Sastri Kota, Lockheed Martin Telecommunications

Pradeep Samudra, Samsung Telecom America, Inc.

Raj Jain is now at Washington University in Saint Louis,
jain@cse.wustl.edu <http://www.cse.wustl.edu/~jain/>



- ❑ Guaranteed Frame Rate
- ❑ Goals
- ❑ Options: Tagging, Buffer Management, Queuing
- ❑ Simulation Results
- ❑ Summary
- ❑ Recommendations

Guaranteed Frame Rate (GFR)

- ❑ Minimum rate guarantee for frames
- ❑ Fair share of unused capacity
- ❑ GCRA like conformance definition
- ❑ Two proposed methods:
 - ❑ FIFO queuing with tagging
 - ❑ Per-VC queuing with per-VC buffer management


GFR (Cont)

- ❑ In April meeting it was shown
 - ❑ Difficult to do GFR for TCP traffic with FIFO queuing and tagging
 - ❑ Can do GFR with per-VC queuing and tagging
- ❑ Per-VC based buffer management was not studied

Goals

- ❑ Explore three options for providing GFR
 - ❑ Tagging (policing)
 - ❑ Buffer Management
 - ❑ Queuing
- ❑ Compare network based tagging vs end system tagging?
- ❑ Compare MCR guarantee to CLP0 vs MCR guarantee CLP0+1?

GFR Options

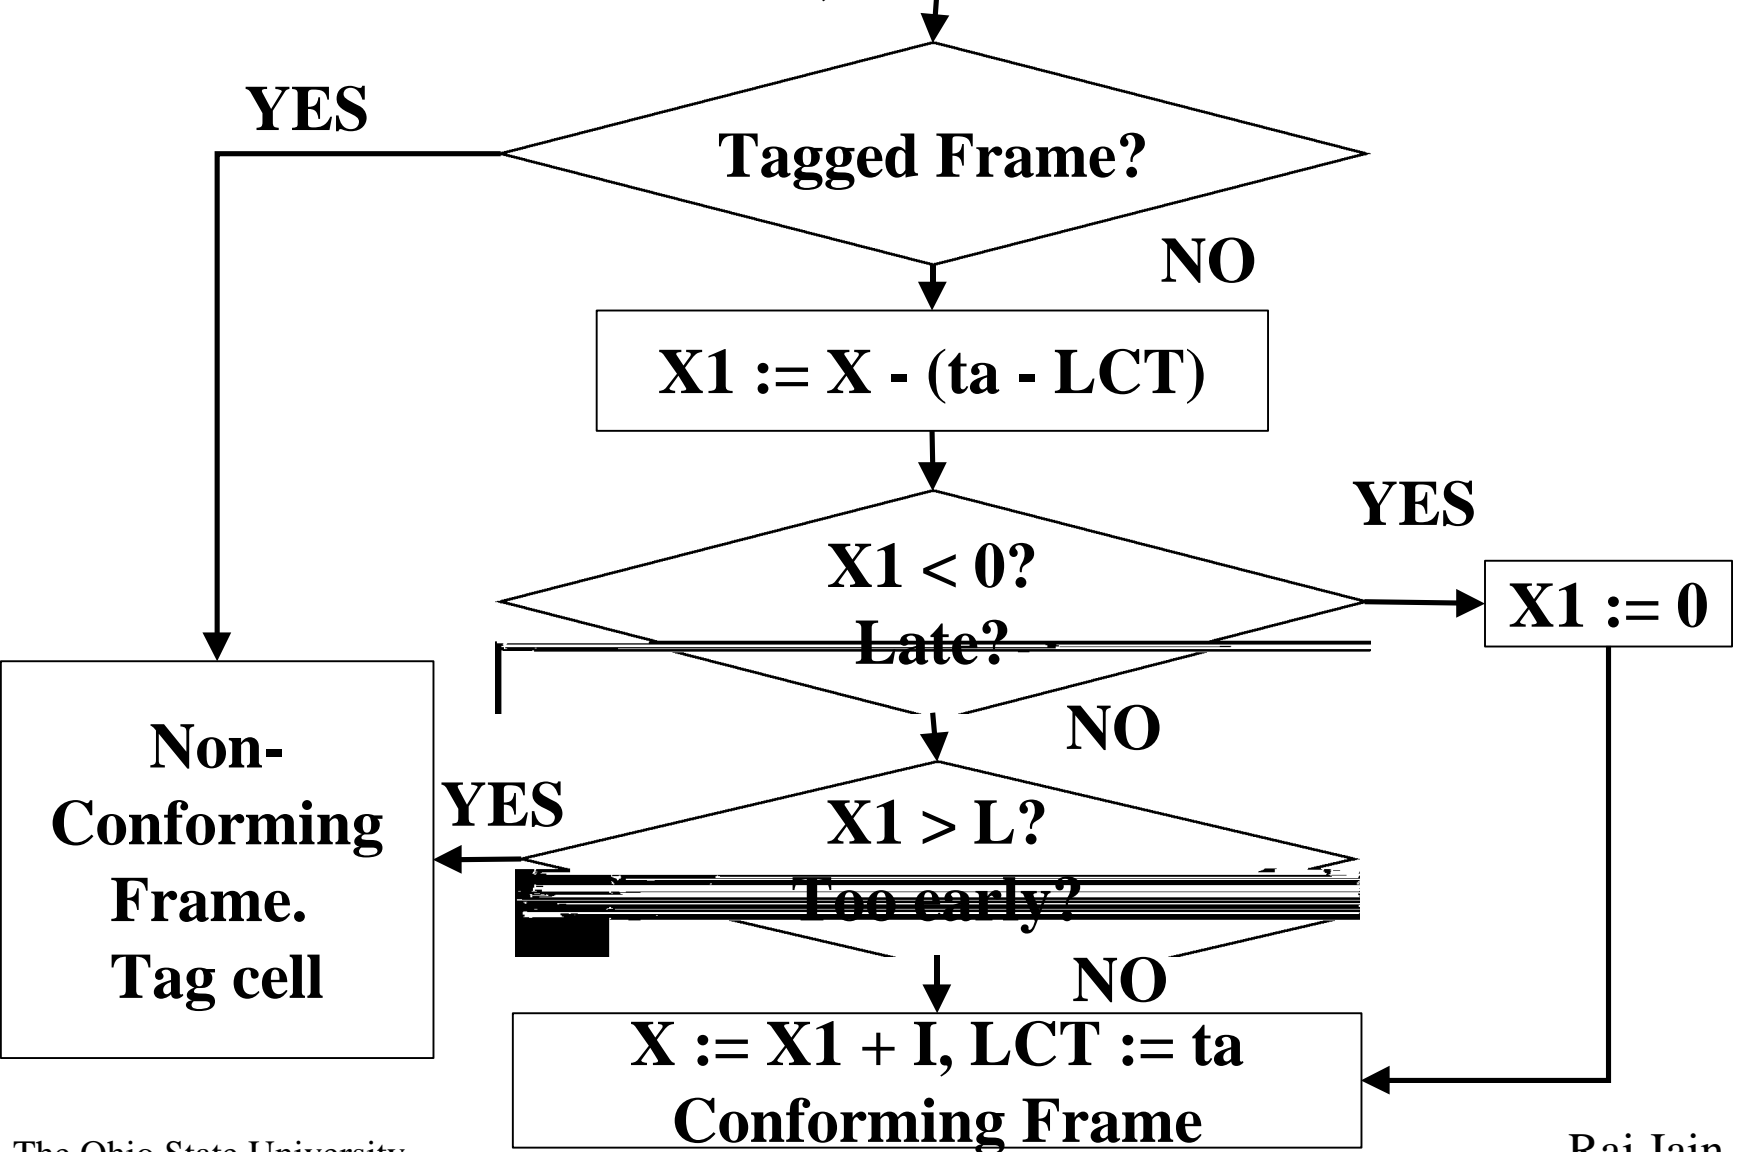


Queuing	Per-VC	FIFO
Buffer Management	Per-VC Thresholds	Global Threshold
Tag-sensitive Buffer Mgmt	2 Thresholds	1 Threshold


Tagging

- ❑ Network based tagging = Policing
- ❑ Continuous state leaky bucket version of the GFR conformance definition:
 - ❑ MCR = Frame rate in cells/sec
 - ❑ MBS = $2 \times \text{CPCS} - \text{SDU size}$
 - ❑ BT = $(\text{MBS} - 1) / (1/\text{MCR} - 1/\text{PCR})$
 - ❑ LCT = Last Compliance Time
 - ❑ CDVT = Tolerance for MCR
 - ❑ X = Leaky bucket counter (nominal arrival time for next cell)
 - ❑ X1 = Local variable

First cell of frame arrives at time t_a .
 $I = 1/MCR, L = CDVT + BT/2$




Non-first cell of a frame arrives at time t_a .



- ❑ Do not drop the last cell of a frame regardless of CLP state unless you drop the entire frame.

Buffer Management



- ❑ K = Buffer Size (cells)
- ❑ R = Congestion Threshold, X = Buffer Occupancy
- ❑ Y_i = Buffer Occupancy of VC_i
- ❑ L_i = Number of untagged cells of VC_i in buffer
- ❑ W_i = Weight of VC_i (based on MCR)
- ❑ N_a = Number of active VCs
- ❑ Z = Fairness threshold

Weighted Buffer Allocation

- When the first cell of a frame arrives:

IF $(X < R)$ THEN

Accept cell and frame

ELSE IF $(X > R)$ THEN

IF $((L_i < R * W_i)$ AND (Untagged)) THEN

Accept cell and frame

ELSE IF $((Y_i - R * W_i) N_a < Z(X - R))$ THEN

Accept cell and frame

ELSE Drop cell and frame


Buffer Management (Cont)

- ❑ Per-VC buffer management controls the entry of frames into the switch buffers.
- ❑ In the absence of network based tagging and per-VC buffer management, VCs that send excess untagged traffic do better than those that tag all their non-conforming traffic
⇒ Per-VC buffer management is needed in the absence of network based tagging

Queuing

- ❑ FIFO versus Per-VC queuing
- ❑ We implemented a WFQ like scheduling policy

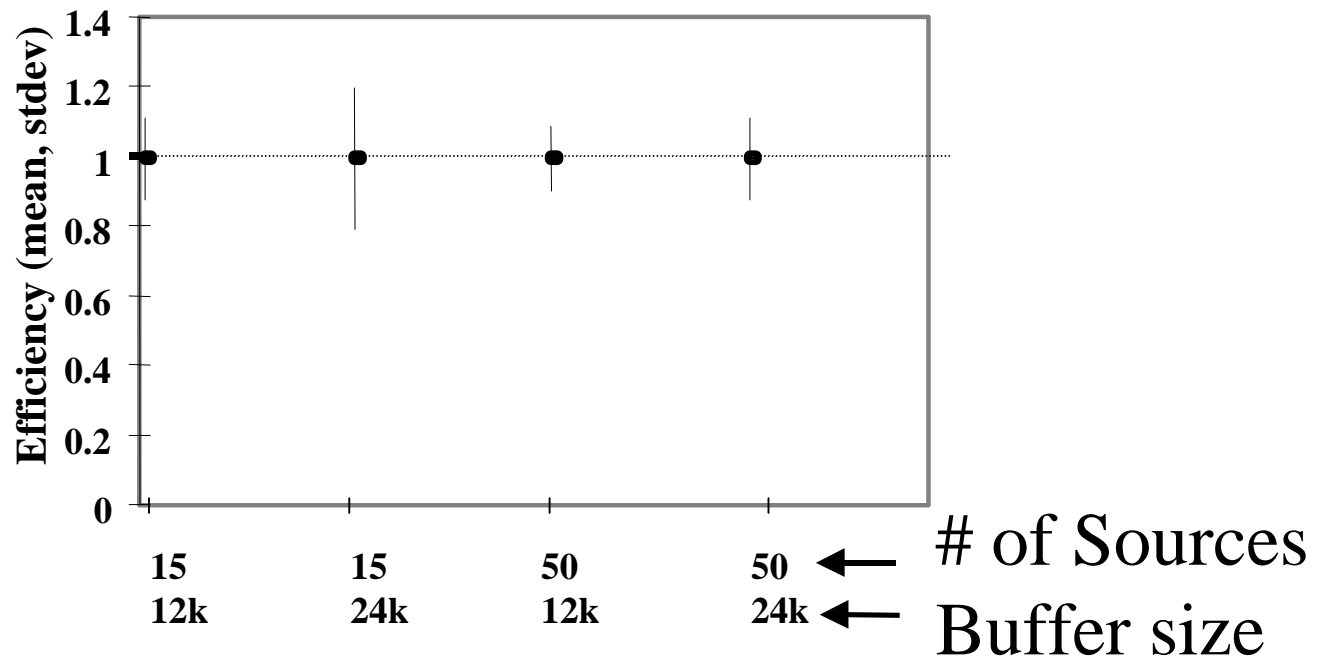
Simulation Experiment



← x Km → | ← x Km → | ← x Km → |


- ❑ N identical infinite TCP sources
- ❑ Link Delay: 5 ms.
- ❑ Link Capacity = PCR = 155.52 Mbps (147.9 Mbps after SONET overhead)
- ❑ Tried both equal and unequal MCR allocations to TCP sources

Equal Rate Allocations




- ❑ Used only per-VC buffer management (sel. drop) with FIFO queuing
- ❑ Bars = standard deviation. Large bars \Rightarrow Unfairness
- ❑ **May allocate equal rates for symmetrical TCP sources with per-VC buffer management**

Unequal Rate Allocations



- ❑ Used per-VC tag sensitive buffer management (WBA) with FIFO queuing
- ❑ Number of sources : 15.
- ❑ 5 Groups with rates = 2.6, 5.3, 8, 10.7, 13.5 Mbps
- ❑ **Cannot allocate unequal rates with FIFO queuing**

Unequal Rate Alloc (Cont)



- ❑ Used only per-VC queuing/scheduling and a single global EPD threshold (not tag sensitive)
- ❑ Number of sources : 15.
- ❑ 5 Groups with MCR = 2.6, 5.3, 8, 10.7, 13.5 Mbps
- ❑ **Can allocate unequal rates with per-VC queuing**


The Role of Tagging

- ❑ End system tagging:
 - ❑ Semantic priority for untagged frames
 - ❑ CLP0 stream has meaning for the end to end performance
- ❑ Network Based tagging:
 - ❑ Conformance of frames
 - ❑ CLP0 stream does not have any special meaning for the end to end performance
- ❑ Network may tag all frames of some VCs to indicate low priority VCs.

Tagging (Cont)

- ❑ Per-VC queuing is needed to make per-VC MCR guarantees
- ❑ FBA + scheduling is needed for fair allocation of excess bandwidth.
- ❑ If guarantees are made to CLP0+1 stream THEN Per-VC queuing + scheduling + FBA is sufficient
- ❑ If guarantees are made to the CLP0 stream THEN Per-VC tag sensitive buffer management is necessary
- ❑ CLP0 may not have any “meaning” if the network performs tagging

Summary



- Fair and proper scheduling is necessary for fair allocation of excess bandwidth
 Per-VC queuing and scheduling is necessary for per-VC MCR. (FIFO + anything cannot do)
- One global threshold is sufficient for CLP0+1 guarantees
 Two thresholds are necessary for CLP0 guarantees