RPR Transit Buffer Schemes

Necdet Uzun and Pinar Yilmaz

AuroraNetics, Inc.

May 14, 2001

Contents

- Objectives
- What are cut-through and store-and-forward?
- Single/dual priority transit buffer ring contention resolution algorithms
- Simulation results
- Conclusion

Objectives

- Guaranteed **END-TO-END** delay and jitter bound for high priority traffic
- Priority discrimination and separation
  - LP traffic does not affect performance of HP traffic
- No packet loss on the ring
- Maximum available ring throughput
  - Delay/jitter performance of HP traffic is not affected by over provisioning of LP traffic
- Best possible delay and jitter for low priority traffic

Applications

- TDM emulation
- Voice over IP
- Interactive video
- Video streaming
- Web browsing

Delay and jitter requirements

Speech Category	End to End Delay[1]	Peak Jitter
Perfect	100 msec	0 msec.
Good	100 msec	75 msec.
Medium	150 msec	125 msec.
Poor	400 msec	225 msec.

[1] including encoding/decoding delays

Source: ITU and ETSI TIPHON (European Telecommunications Standards Institute "Telecommunications and Internet Protocol Harmonization Over Networks")

What is cut-through?

Two interpretations:

- Cutting through traffic
  - Ring traffic “cuts through” host traffic
  - Versus allowing host traffic (high priority) to be transmitted before low priority ring traffic
- Cut-through buffers
  - No need to store complete packet before starting transmission
  - Versus Store-and-forward, storing complete packet (checking for errors, etc.) before starting transmission

Cutting through

- If all transit traffic is “cutting through” transmit traffic:
  - High-priority transmit packets have to wait for low priority transit packets (incurring more jitter for high priority)
  - Transit packets can be dequeued in “cut-through” or “store-and-forward” fashion
    - Cut-through buffers with cutting through traffic (CT)
    - Store-and-forward buffers with cutting through traffic

Are cut-through buffers needed?

AuroraNetics, Inc.

- High priority traffic demonstrates CBR characteristics: short packet size, does not vary much, small transmission delay
- 1.5KB packet at 1Gbps = 12µsec
- Cut-through buffer is not justified for the link rates considered in RPR (1Gbps and up)

Cut-through (CT) simulation model

- Transit traffic cuts through the transmit traffic
  - Transit packets have priority over transmit packets
  - Transit packets are not stored fully before starting transmission

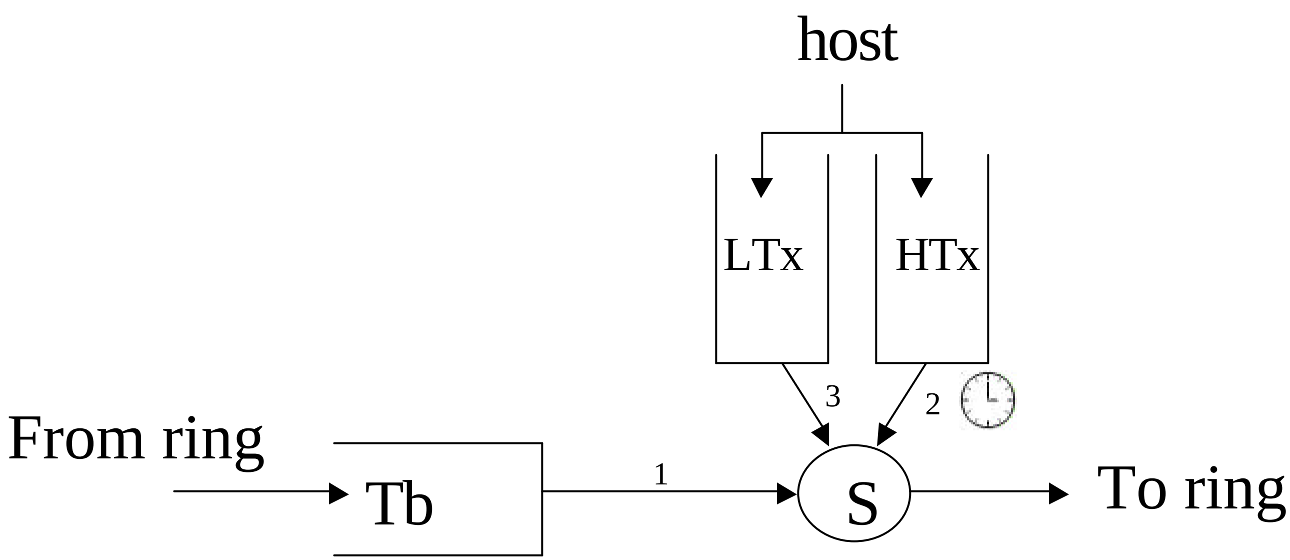

Store and forward (SF) model

- Only high-priority transit traffic cuts through the transmit traffic
  - – Transit packets are fully stored before they are forwarded to the ring (SF)

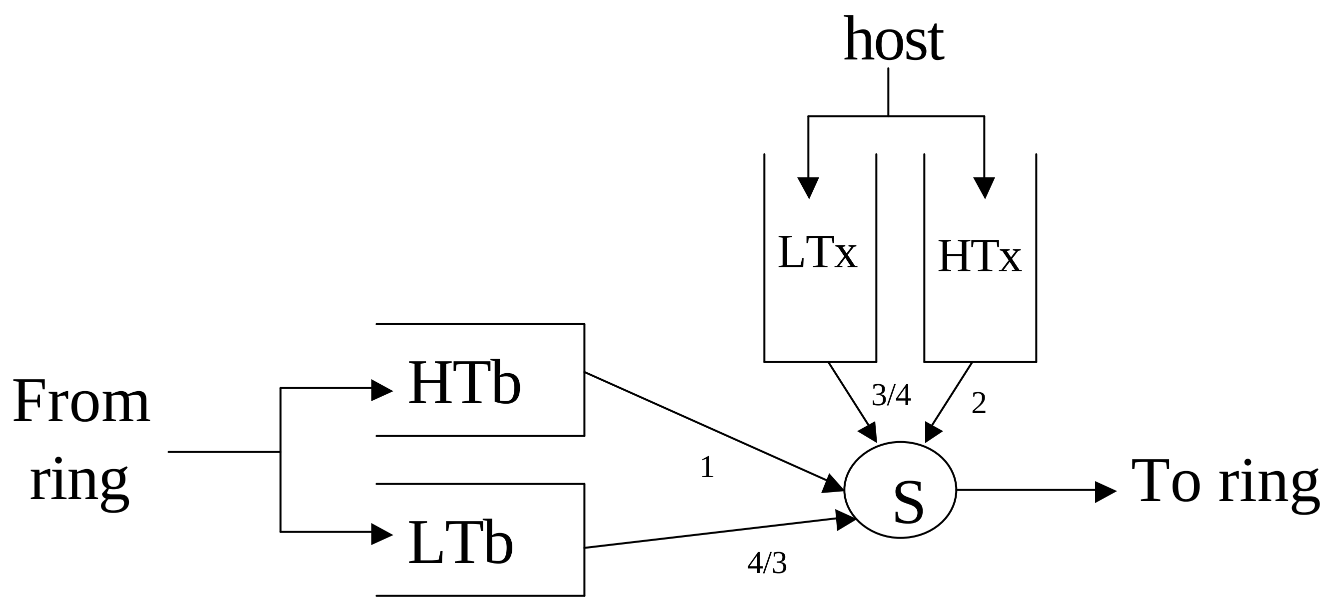

Scenarios

AuroraNetics, Inc.

- 16 node, 100Km OC192 dual ring
- Packet size: 64B(%60), 512B(%20), 1518B(%20)
- Single Tb (CT)
  - Tb = 32KB
  - Cut-through
- Dual Tb (SF)
  - HTb = 32KB
  - LTb = 256KB
  - Store and forward

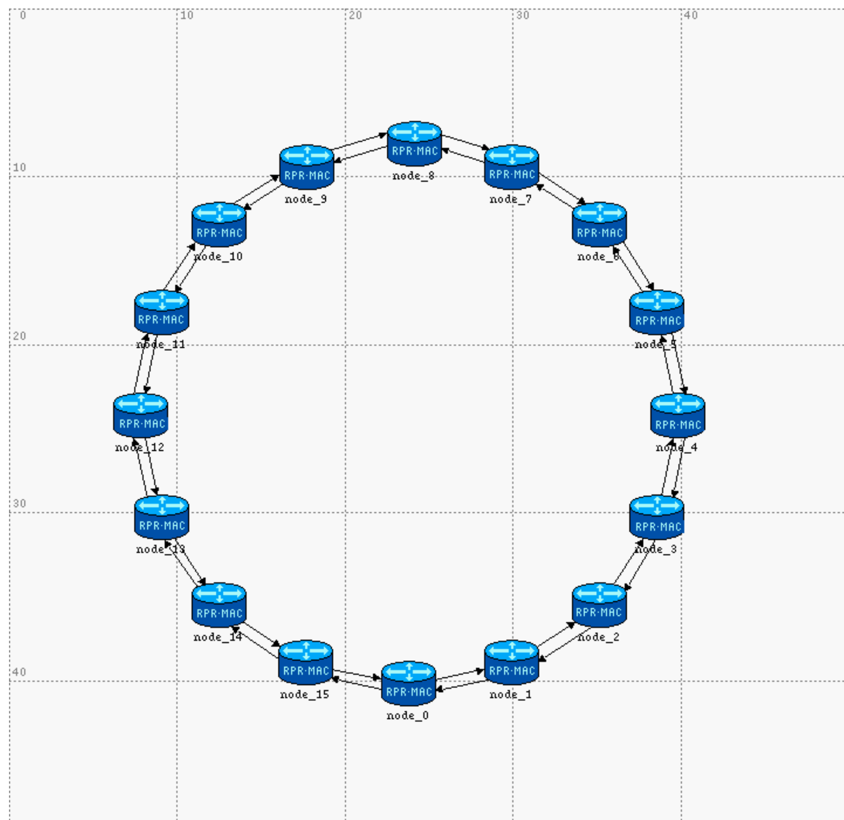

Traffic scenarios: Mesh (Any-to-Any)

AuroraNetics, Inc.

- HTx: 370Mbps CBR
  - 1 tri-modal source per node
- LTx: 2.1Gbps bursts
  - 16 tri-modal sources per node
  - on 1msec, off 9msec, exponential distribution
  - total of ~3.4Gbps LTx per node
- Total traffic injected: ~60Gbps
  - Total HP traffic is ~6Gbps

Traffic scenarios: Hub (Any-to-Hub, Hub-to-Any)

- Node to Hub (node 0):
  - HTx: 430Mbps CBR
    - 1 tri-modal source per node
  - LTx: 2Gbps bursts
    - 1 tri-modal source per node
    - on 1msec, off 1msec, exponential distribution
    - total of ~1Gbps LTx per node
- Hub (node 0) to Node:
  - HTx: 6Gbps CBR (tri-modal source)
  - LTx: 2Gbps bursts
    - 14 tri-modal sources
    - on 1msec, off 1msec, exponential distribution
    - ~1Gbps LTx per source, total of ~14Gbps
- Total traffic injected: ~40Gbps
  - Total HP traffic is ~12Gbps

Note: Node 8 does not send or receive traffic

Mesh simulation results: Jitter

AuroraNetics, Inc.

100 km Mesh High Priority ETE Delay Histogram

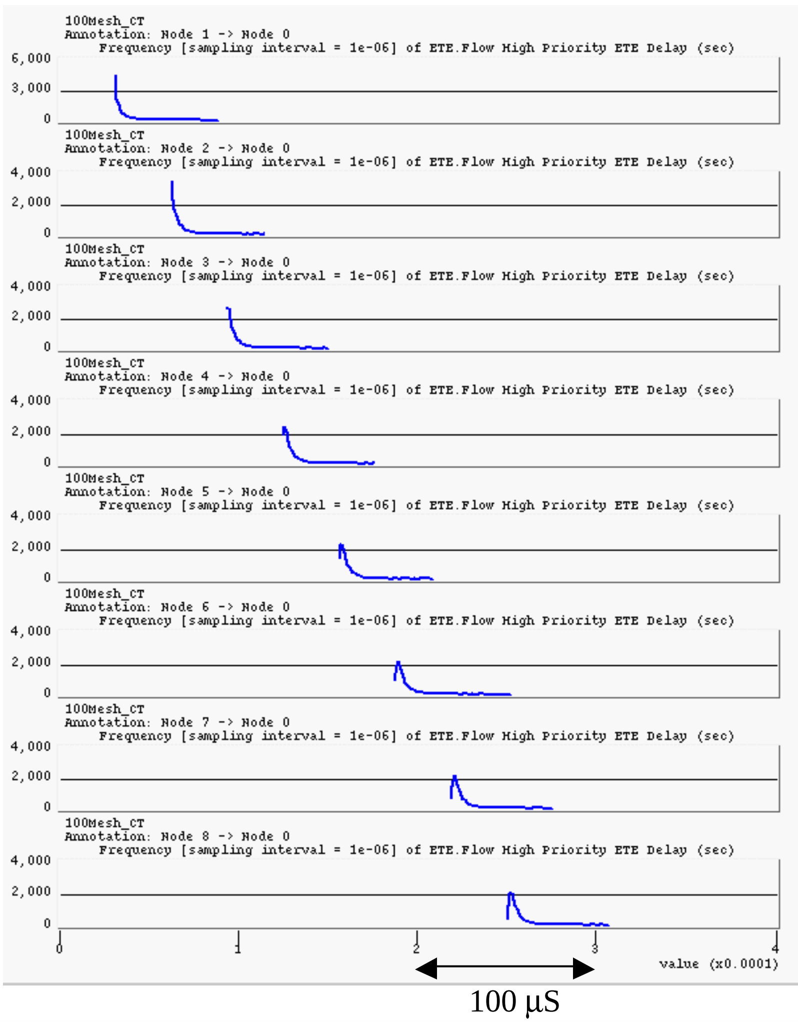

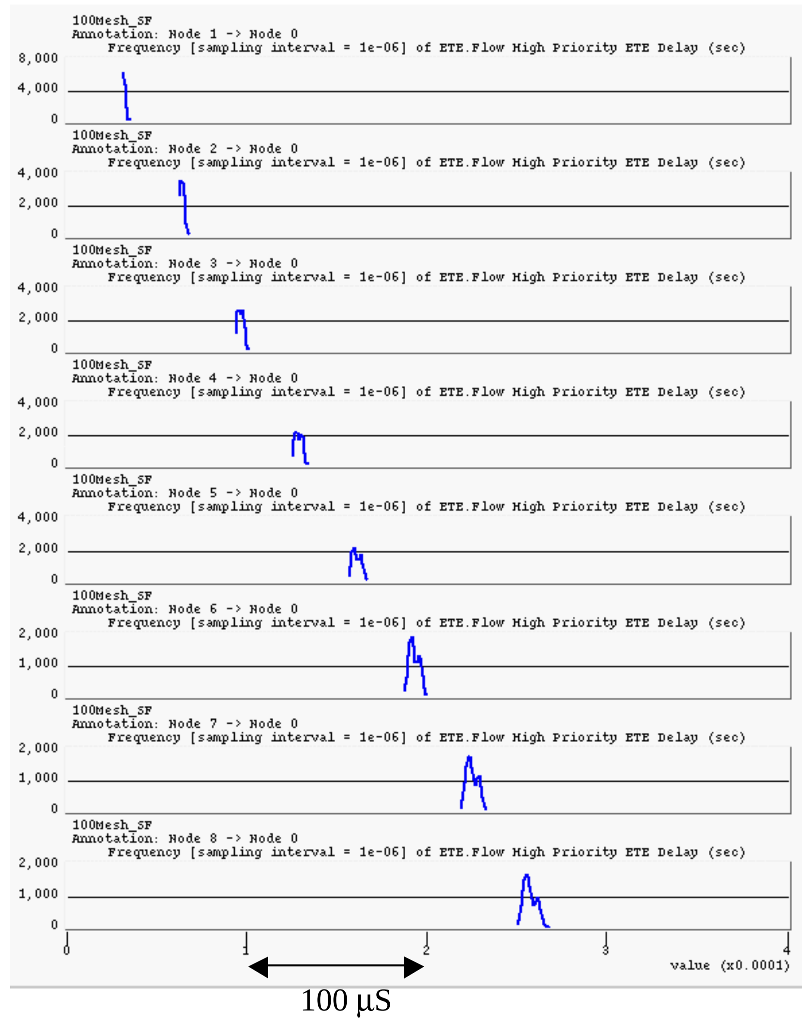

Mesh simulation results: Jitter Close-up

AuroraNetics, Inc.

100 km Mesh High Priority ETE Delay Histogram

Node 1 to Node 0

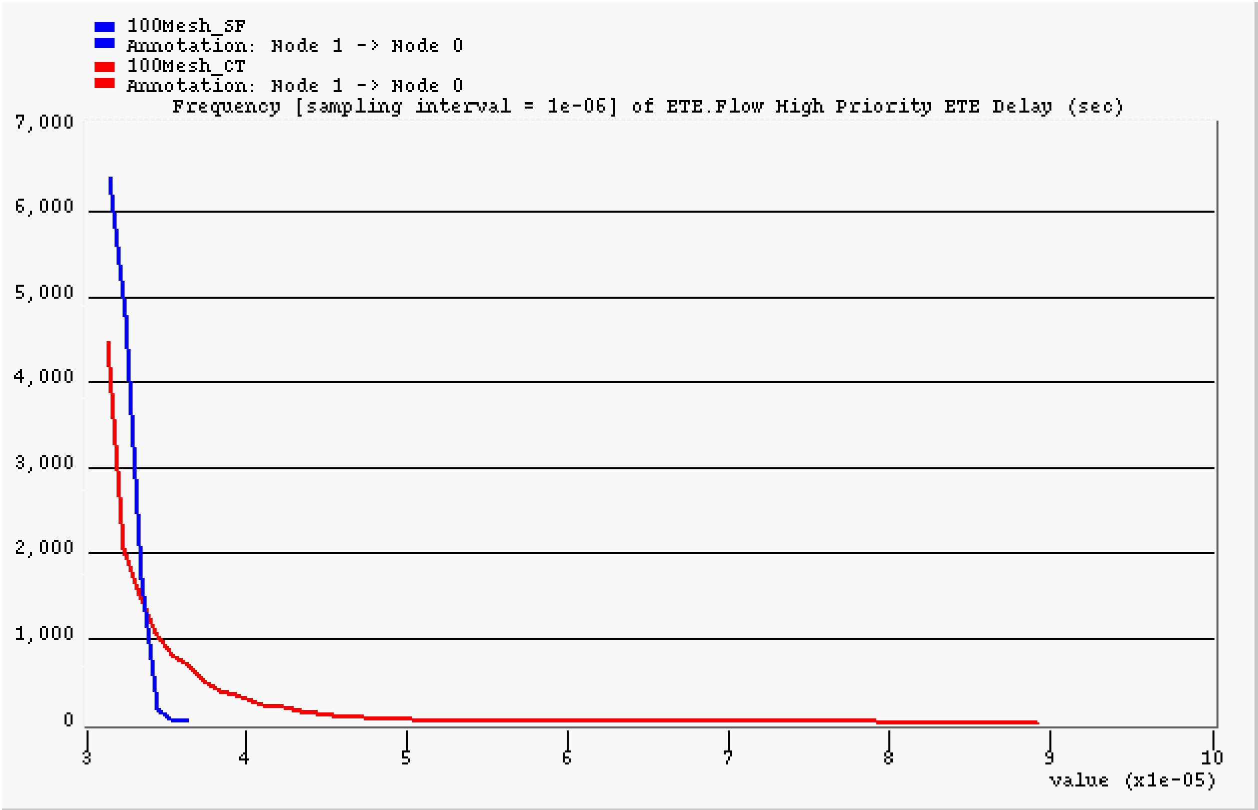

10 µS

Hub simulation results: Jitter

AuroraNetics, Inc.

100 km Hub High Priority ETE Delay Histogram

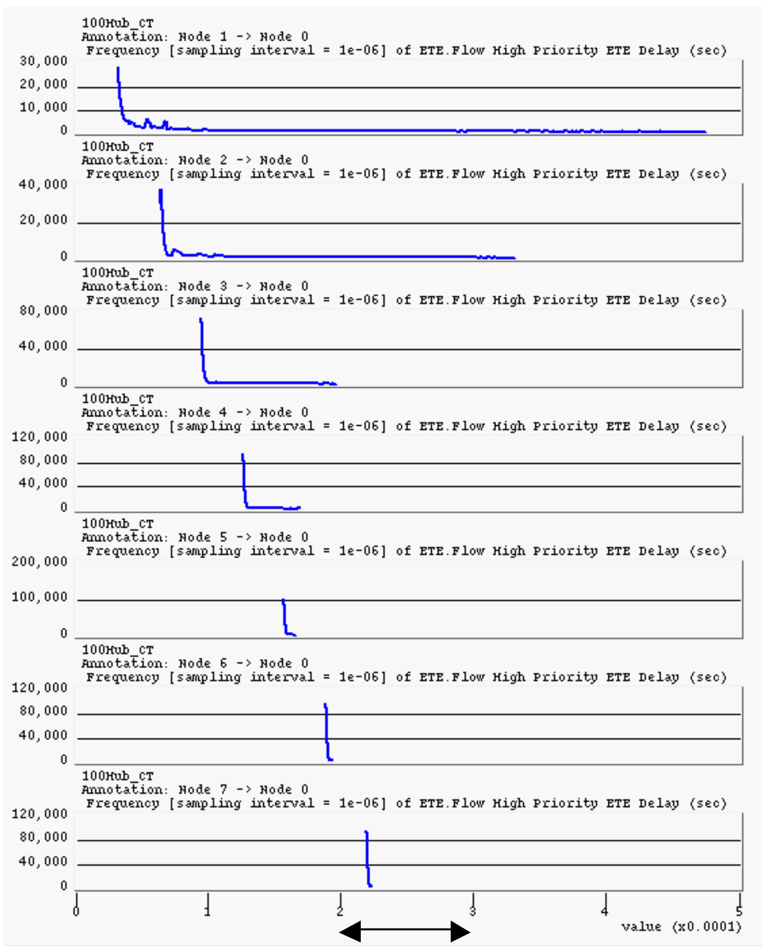

100 µS

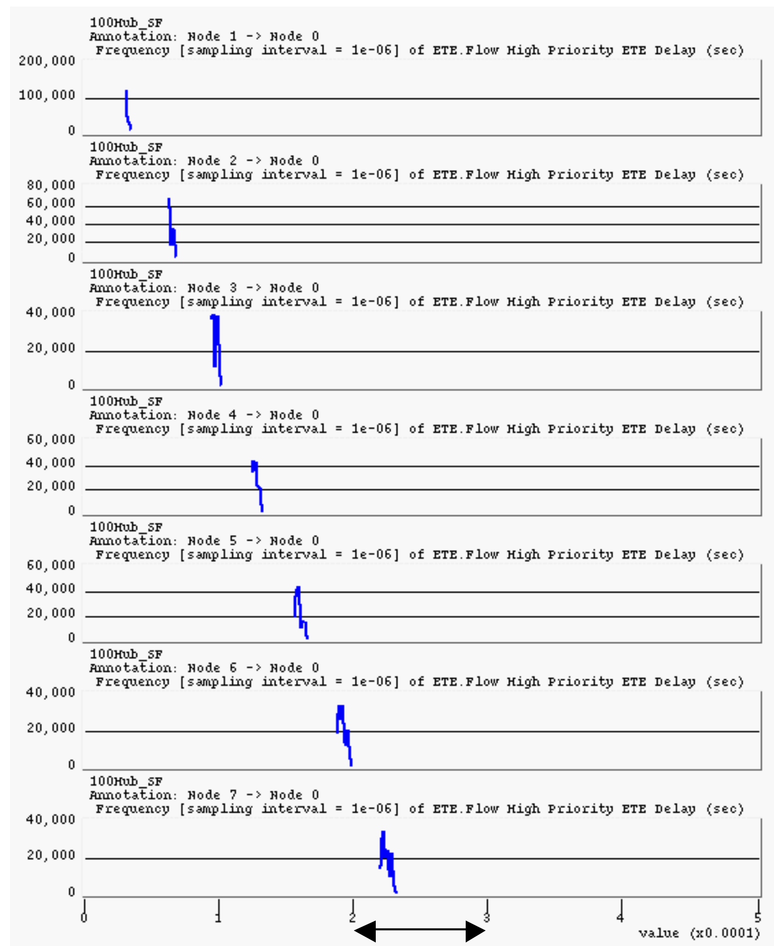

100 µS

Hub simulation results: Jitter Close-up

100 km Hub High Priority ETE Delay Histogram

Node 1 to Node 0

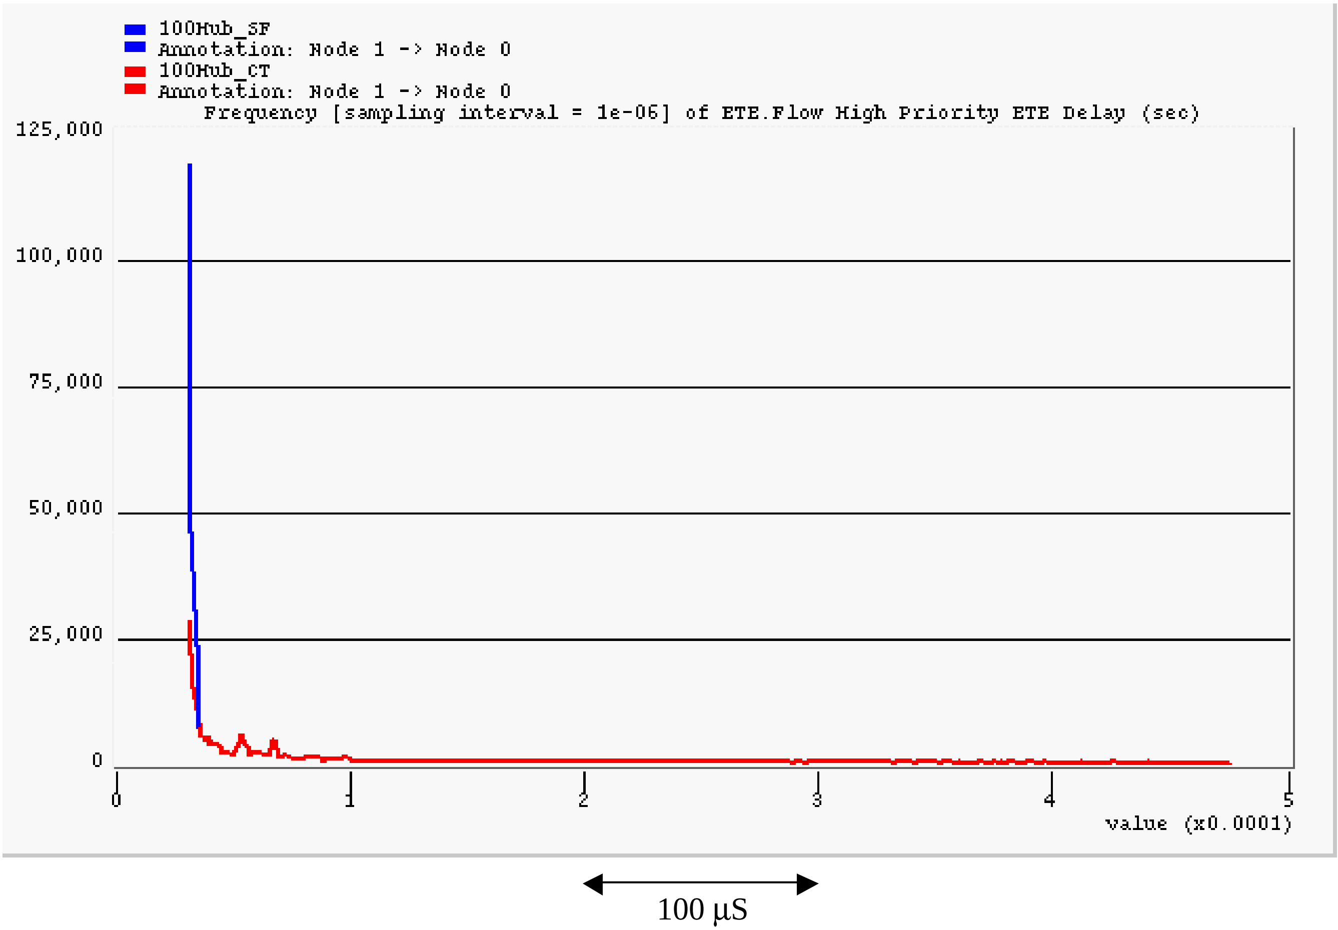

Buffer utilization

Node 1:

CT Tb vs SF HTb

Outer Ring

3KB

2KB

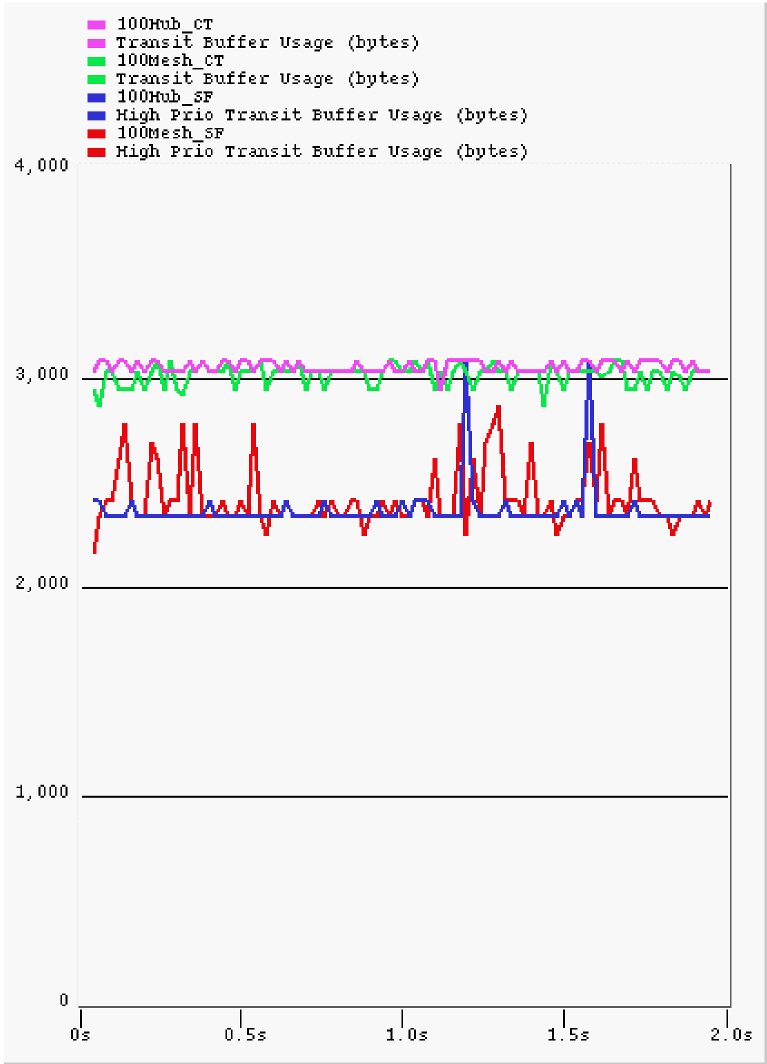

Buffer utilization (cont)

Node 1: SF LTb Outer Ring

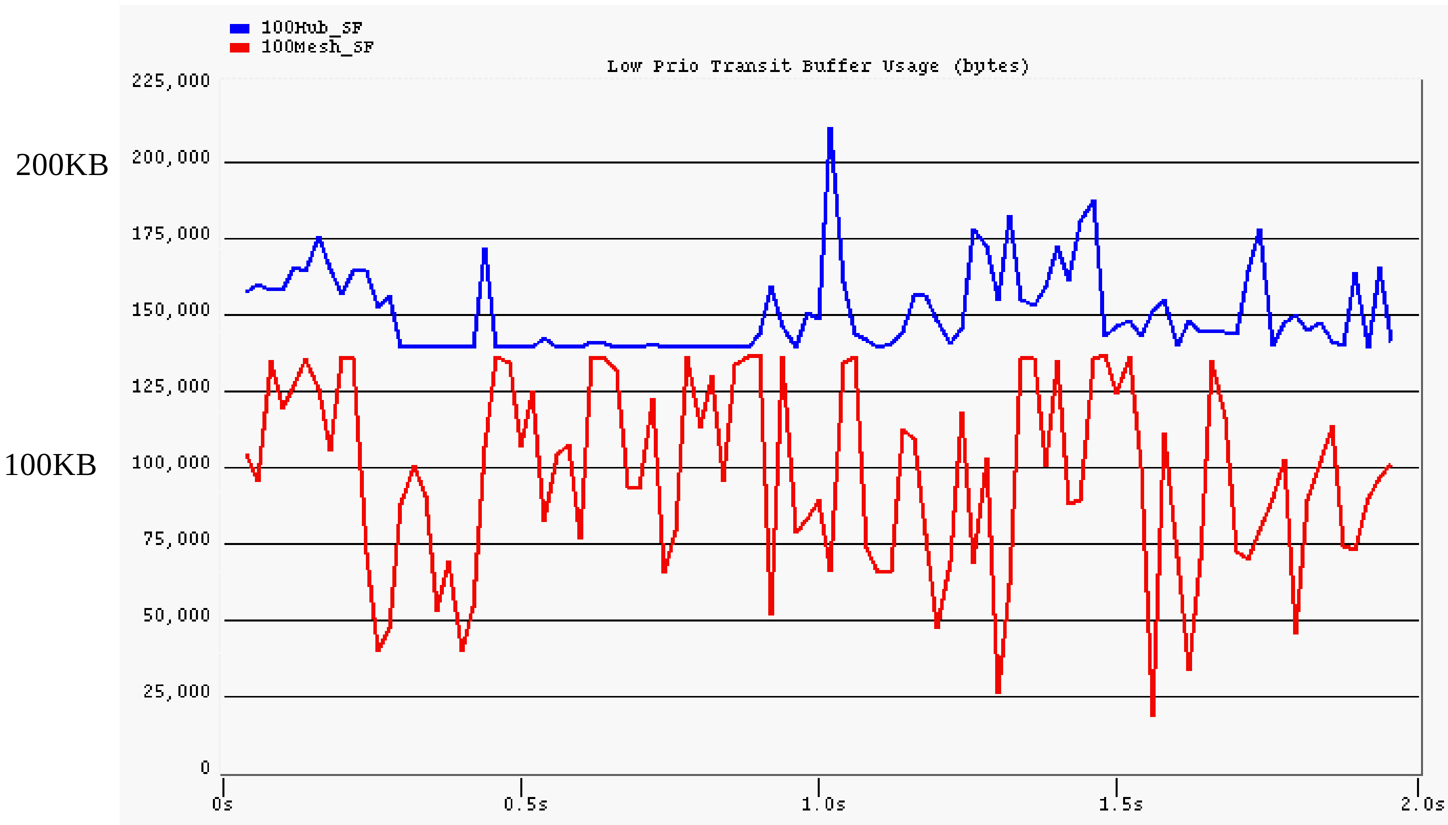

Throughput

AuroraNetics, Inc.

Global Aggregated Traffic Delivered (bps)

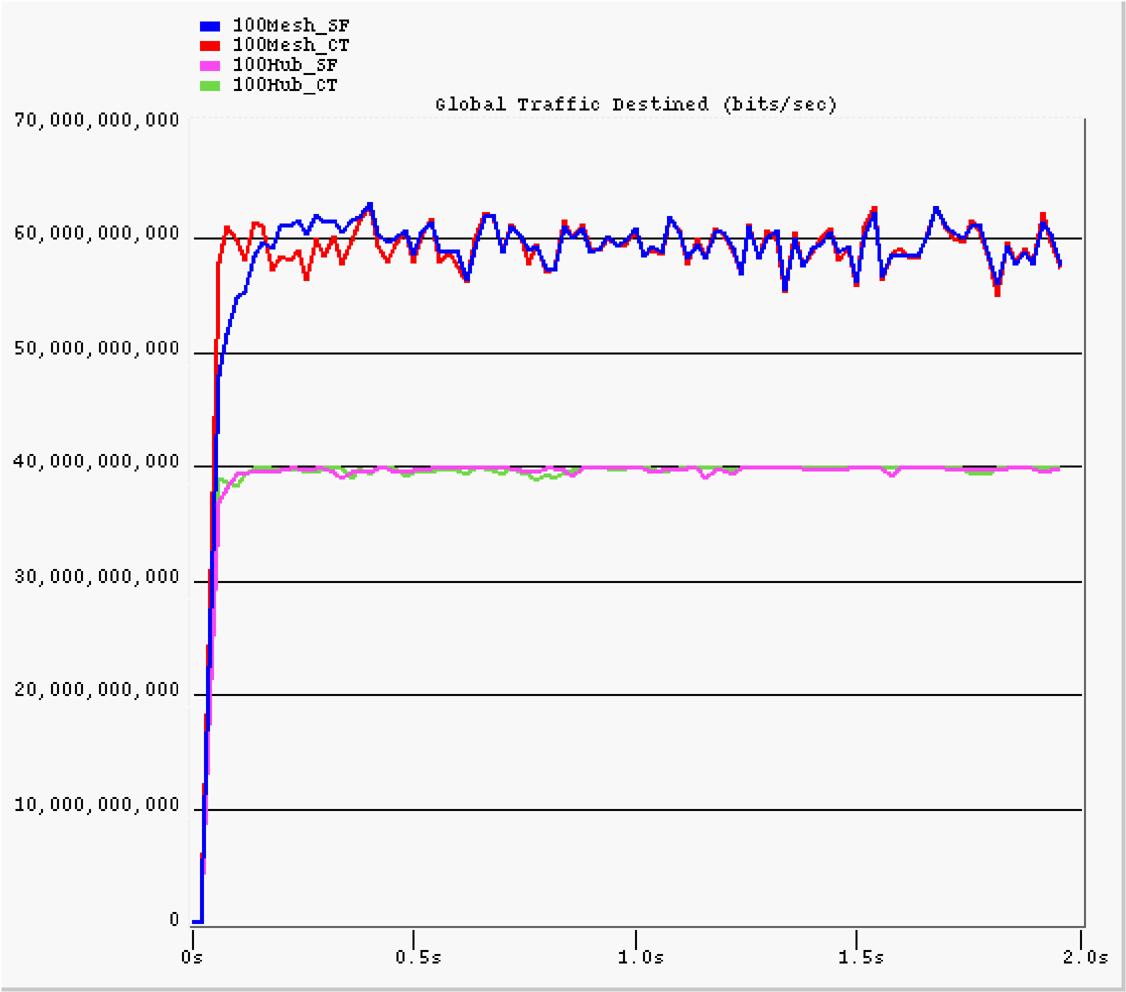

HP jitter comparison

- Cut-through buffers cause more jitter for high priority
  - Transit packets does not give HP transmit packets a chance to get into the ring
    - In SF while a transit packet is being stored, a transmit packet gets a chance to enter to the ring
- Cutting through (single transit buffer) implementation causes unacceptable HP jitter in HUB scenario
  - Node just before the HUB is bombarded with LP transit packets which do not let HP transmit packets to get in to the ring
    - Cut-through buffers make it worse
    - SF buffers may perform better in this implementation
- Multi transit buffer implementation guarantees HP jitter bound by decoupling HP and LP traffic

Summary

– Cutting through transit traffic causes more jitter for high priority

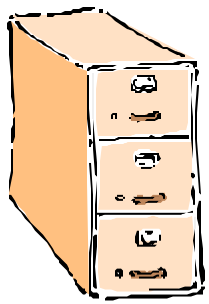

– Store-and-Forward does not require large buffers

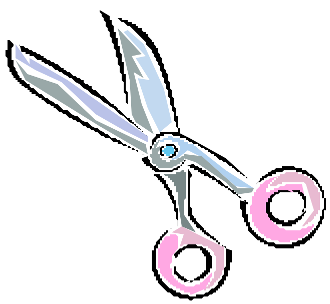

– Multiple transit buffers can result in more throughput (for comparable HP delay jitter) than single transit buffer

AuroraNetics, Inc.

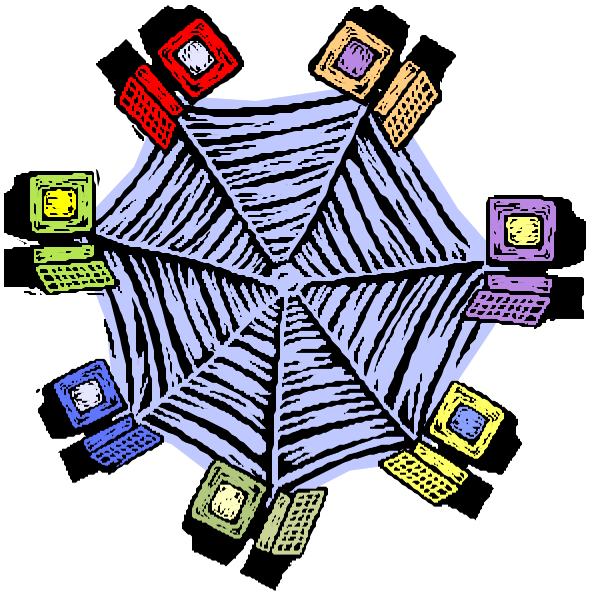

Thank you!

Q & A