SQR

In-network packet loss recovery from link failures for high-reliability datacenter networks

Ting Qu [1, 2] Raj Joshi [2] **Mun Choon Chan**[2]

Ben Leong[2] Deke Guo[1] Zhong Liu[1]

1 國防科學技術大學 National University of Defense Technology

2 NUS National University of Singapore

Data centers around the world

Google's worldwide DC map

Facebook DC interior

Microsoft's DC in Dublin, Ireland

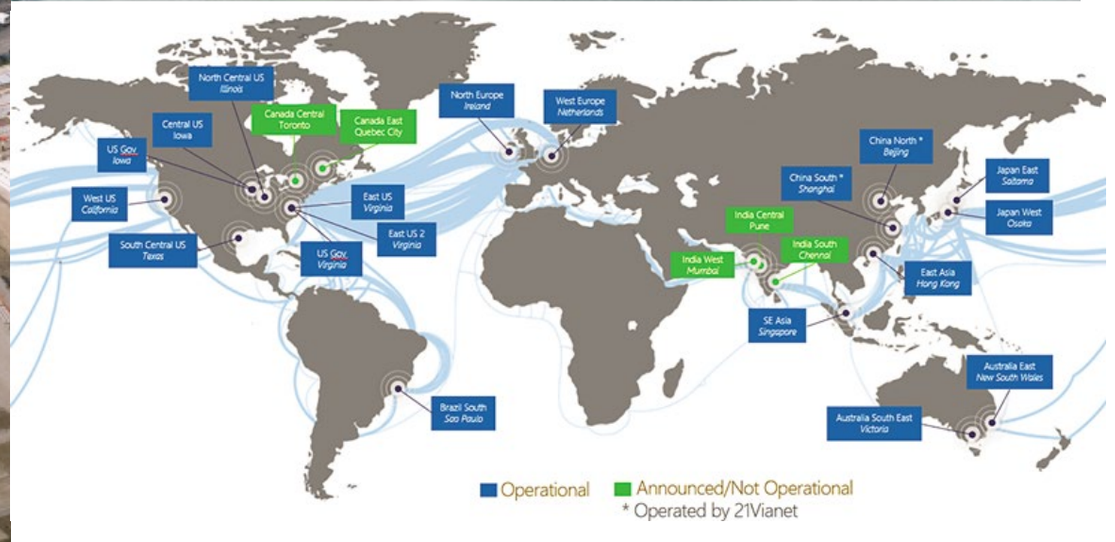

Global Microsoft Azure DC Footprint

Low latency is a key requirement

Web search

e-commerce

database

cache

Low latency for short messages

Better app performance & user experience

Improve Flow Completion Time (FCT)

- DCTCP (sigcomm'10)
- D^3 (sigcomm'11)
- HULL (nsdi'12)
- pFabric (sigcomm'13)
- PASE (sigcomm'14)
- TIMELY (sigcomm'15)
- FUSO (atc'16)
- Homa (sigcomm'18)
- HPCC (sigcomm'19)

…

But very few work specifically address how link failures impact FCT

Link failures are common

- Gill et al. [1] reported:
  - Link failure are common and can cause loss of a large number of small packets.
  - The 95th percentile value of link failure is 136 times per day during their measurement period.

[1] Phillipa Gill, Navendu Jain, and Nachiappan Nagappan. 2011. Understanding network failures in data centers: measurement, analysis, and implications. In Proceedings of SIGCOMM.

Link failure management

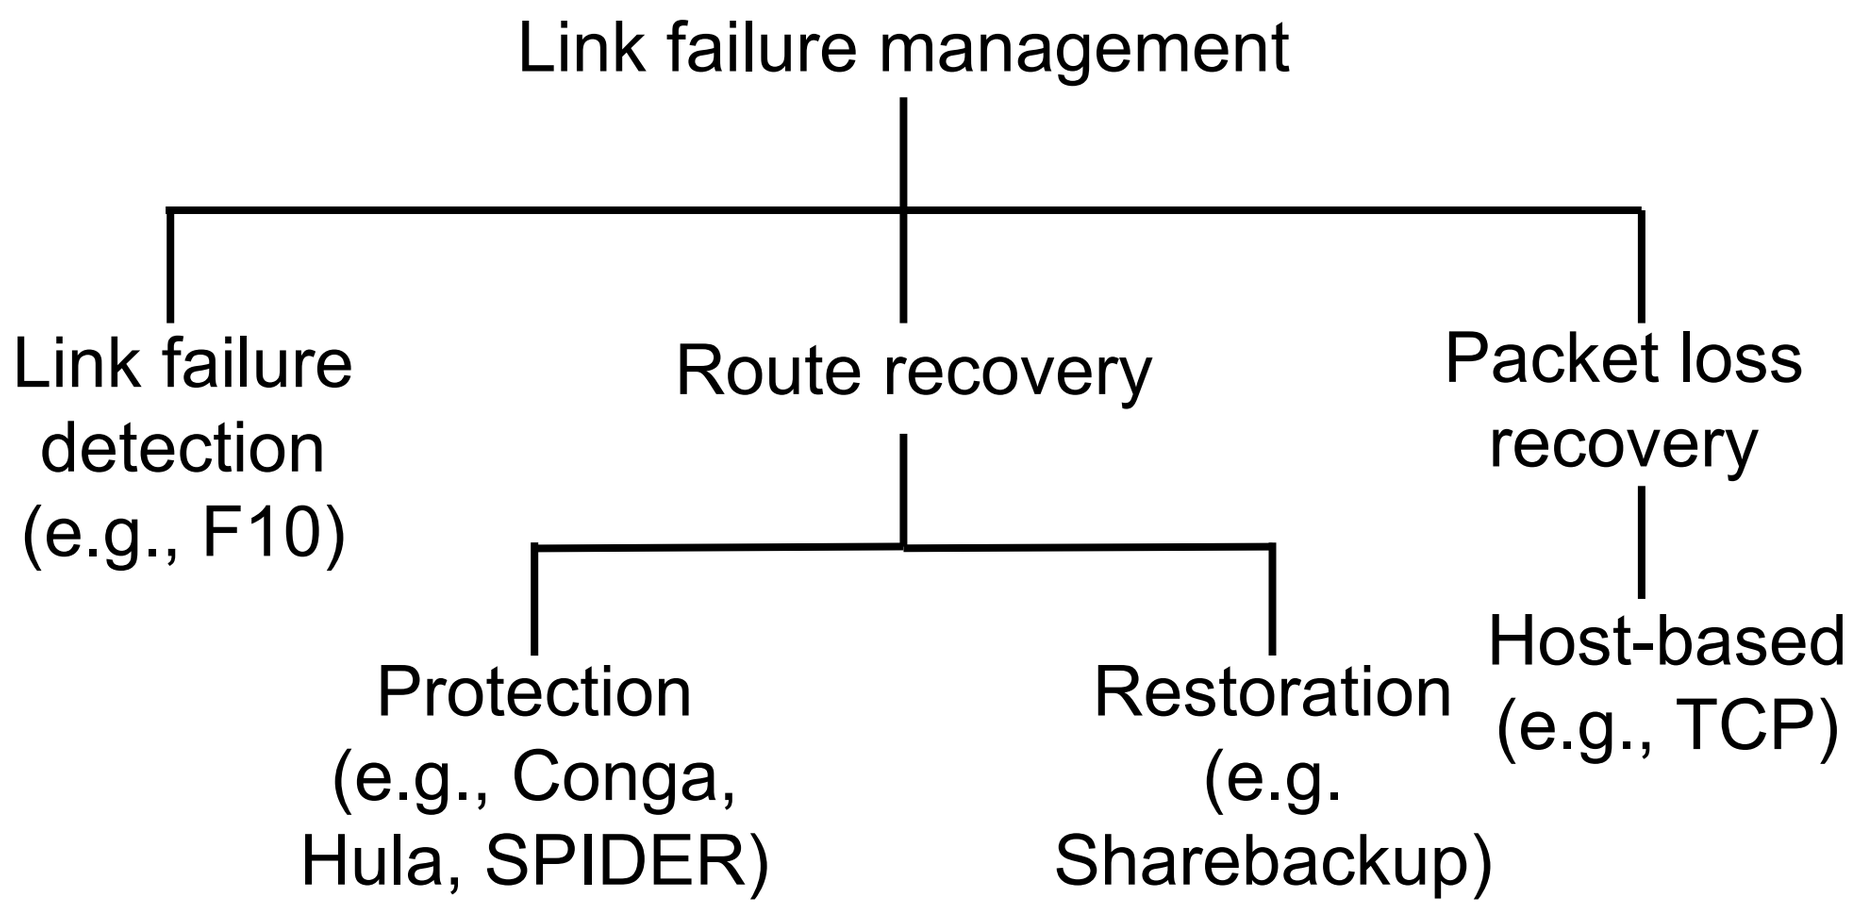

Host-based pkt loss recovery can lead to much longer flow completion time (FCT) for short flows

Link failure case

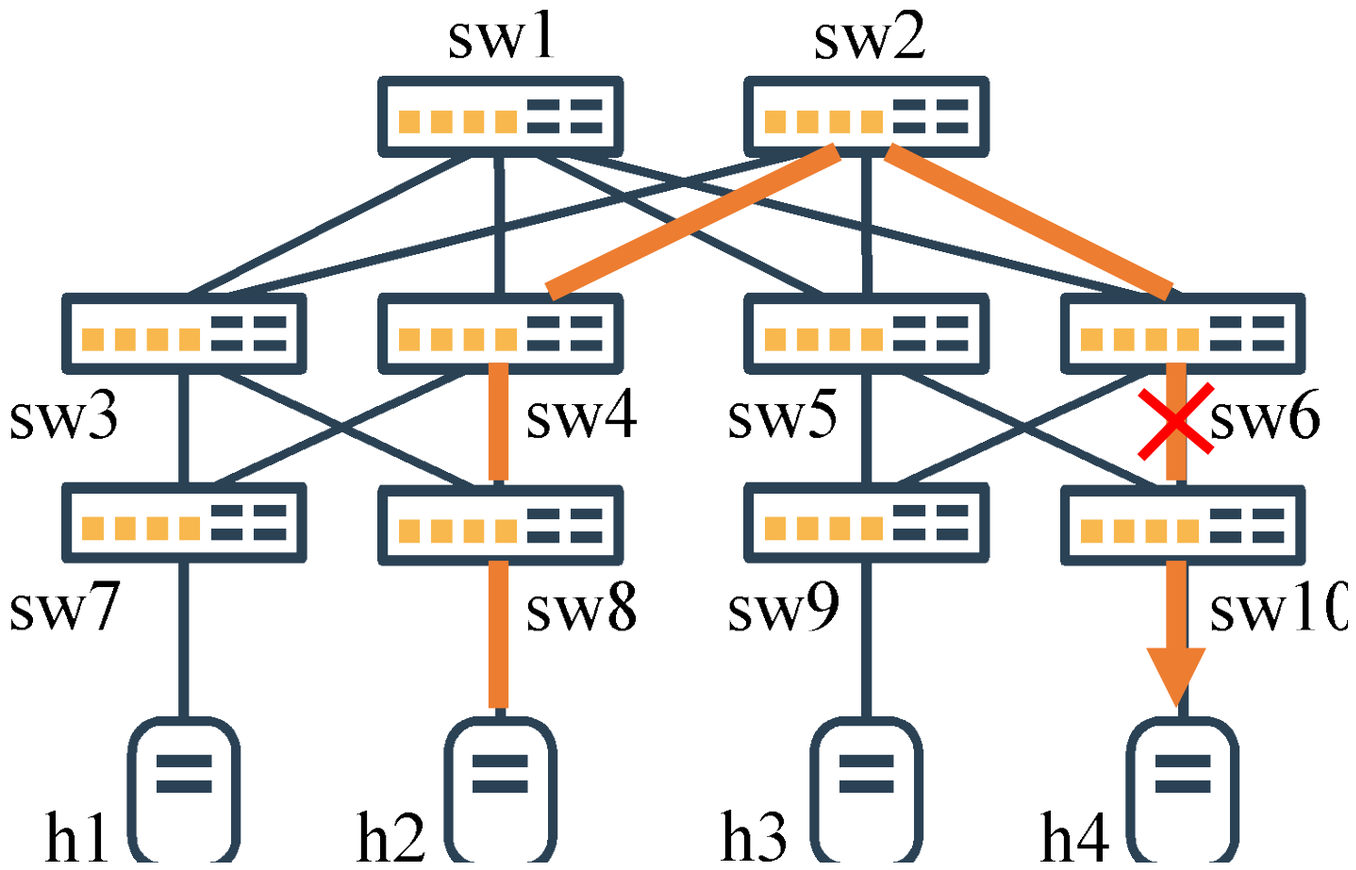

Link detection time
30us
(F10, NSDI'13)

+

route reconfiguration time
730us
(ShareBackup, sigcomm'18)

=

760us

Long FCT under link failure

Host based recovery is a major contributor to the large increase in FCT

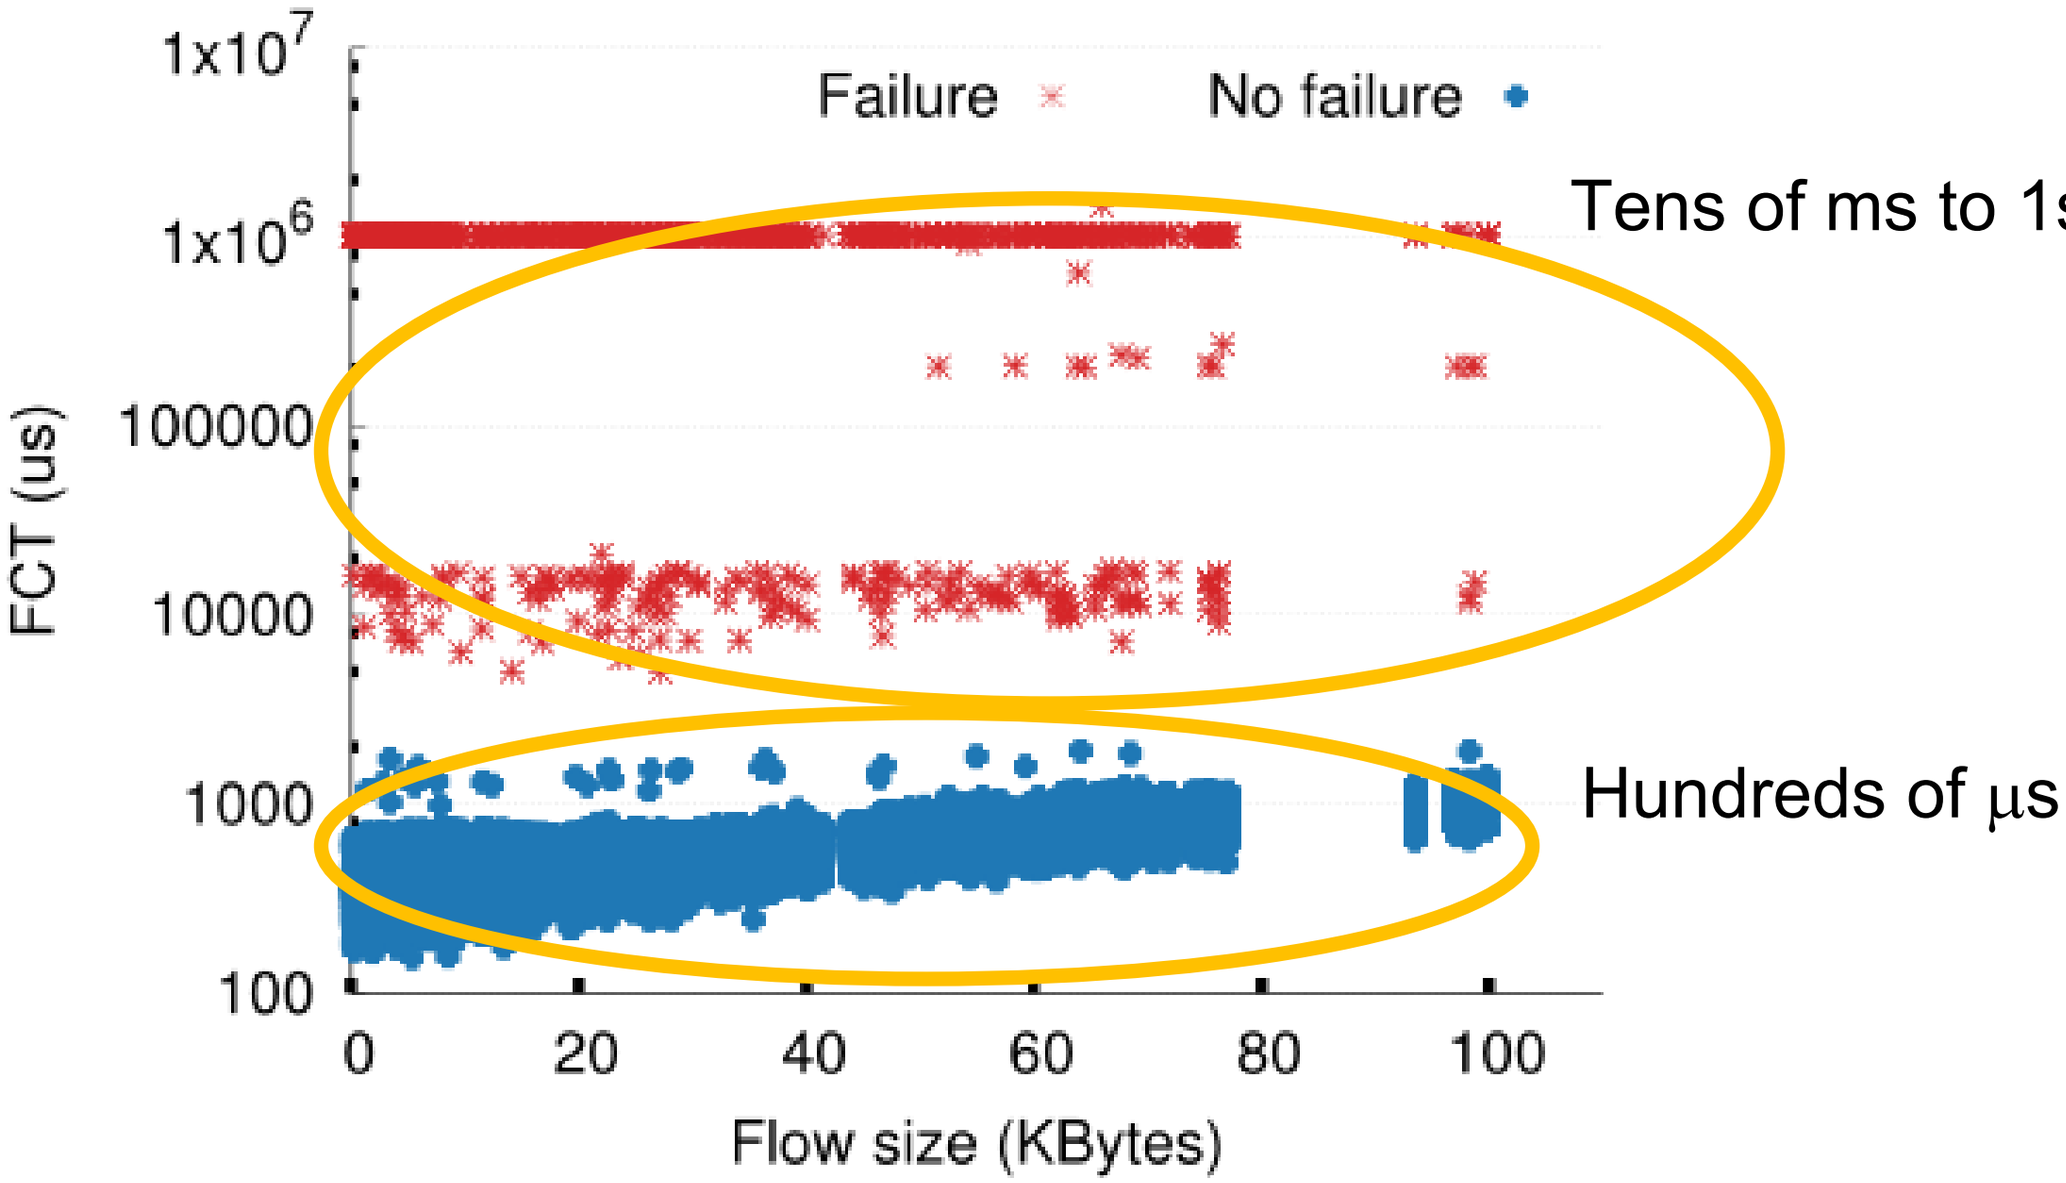

Why does host-based recovery increase FCT significantly?

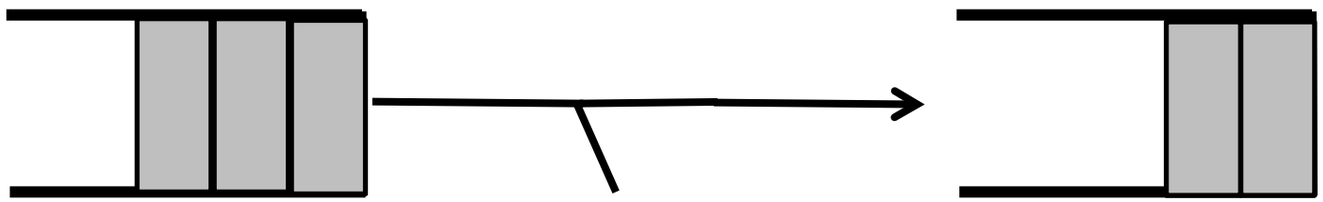

- Packet losses in the TCP three-way handshake
  - Wait at least 1s and retransmit
- Packet losses in the middle of a cwnd
  - Fast retransmission: 1RTT (100s of us)
- Packet losses at the tail of a cwnd
  - Retransmission timeout: several ms

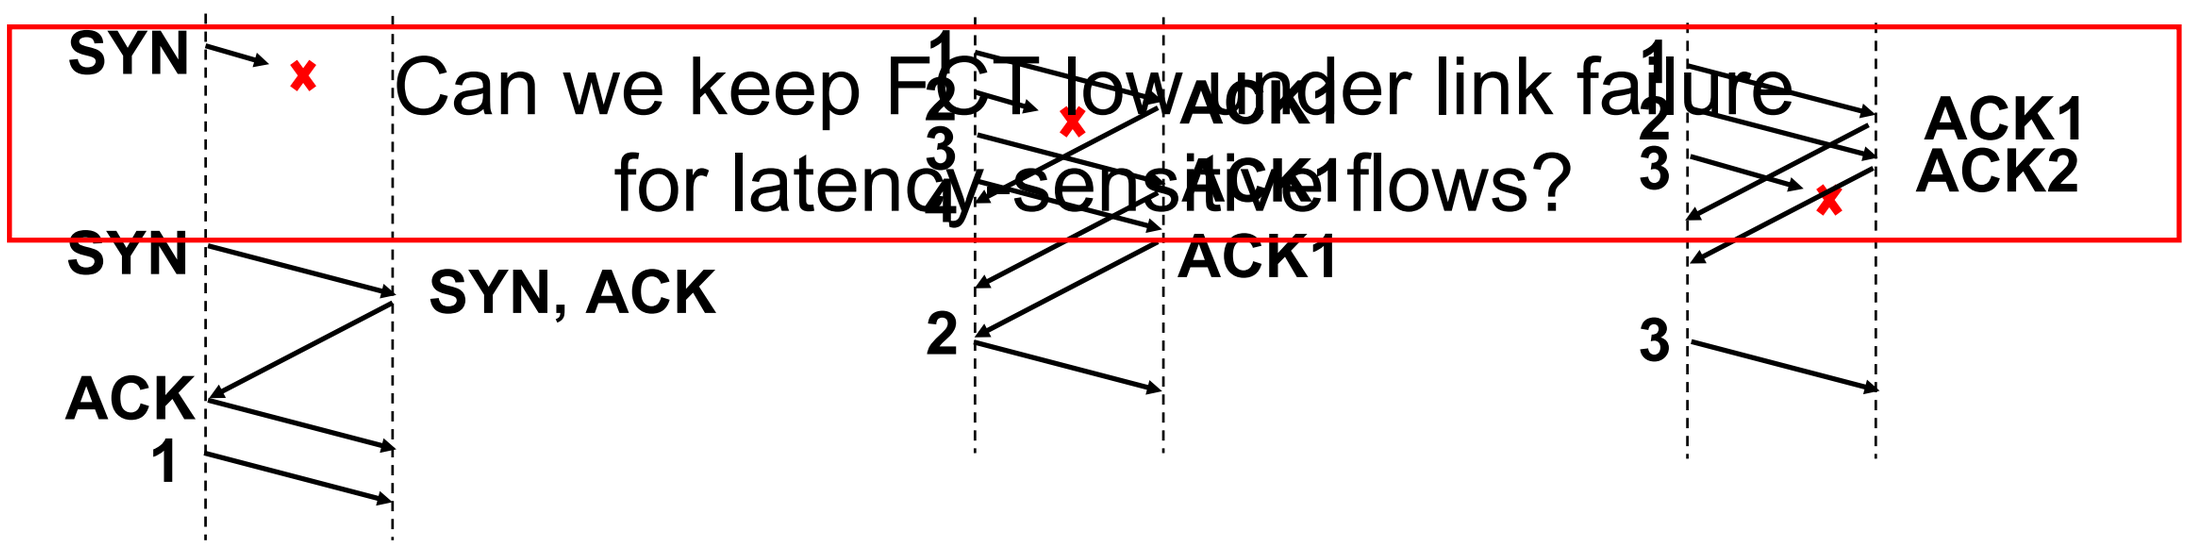

Can we keep FCT low under link failure for latency-sensitive flows?

Our solution: SQR

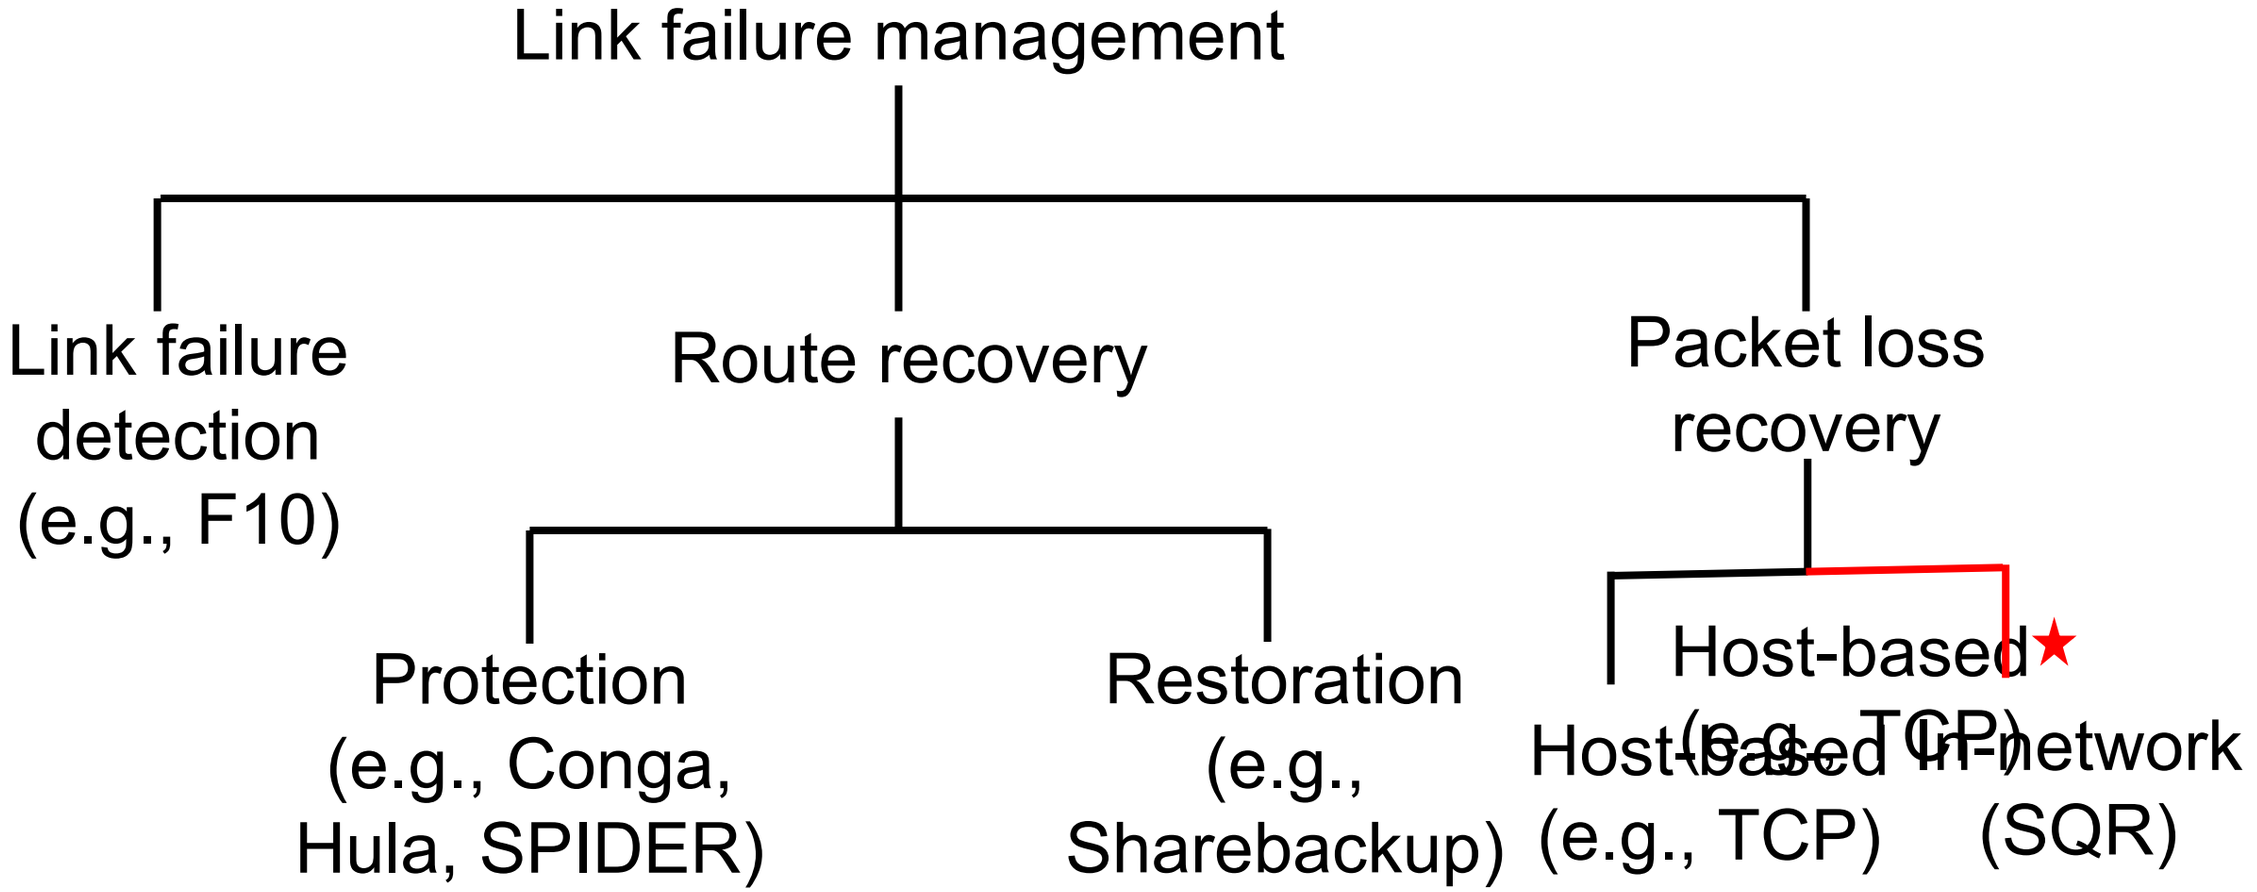

The network is the "right" place to perform packet loss recovery

How does SQR keeps FCT low when there is link failure?

Objective:

- **Mask the effect of packet loss from the end-points** during link failure detection time and route reconfiguration time (route failure time).

Key idea:

- **Continuously cache recently sent packet in the switch for a duration equal to the route failure time**

Is it feasible to cache pkts on switch?

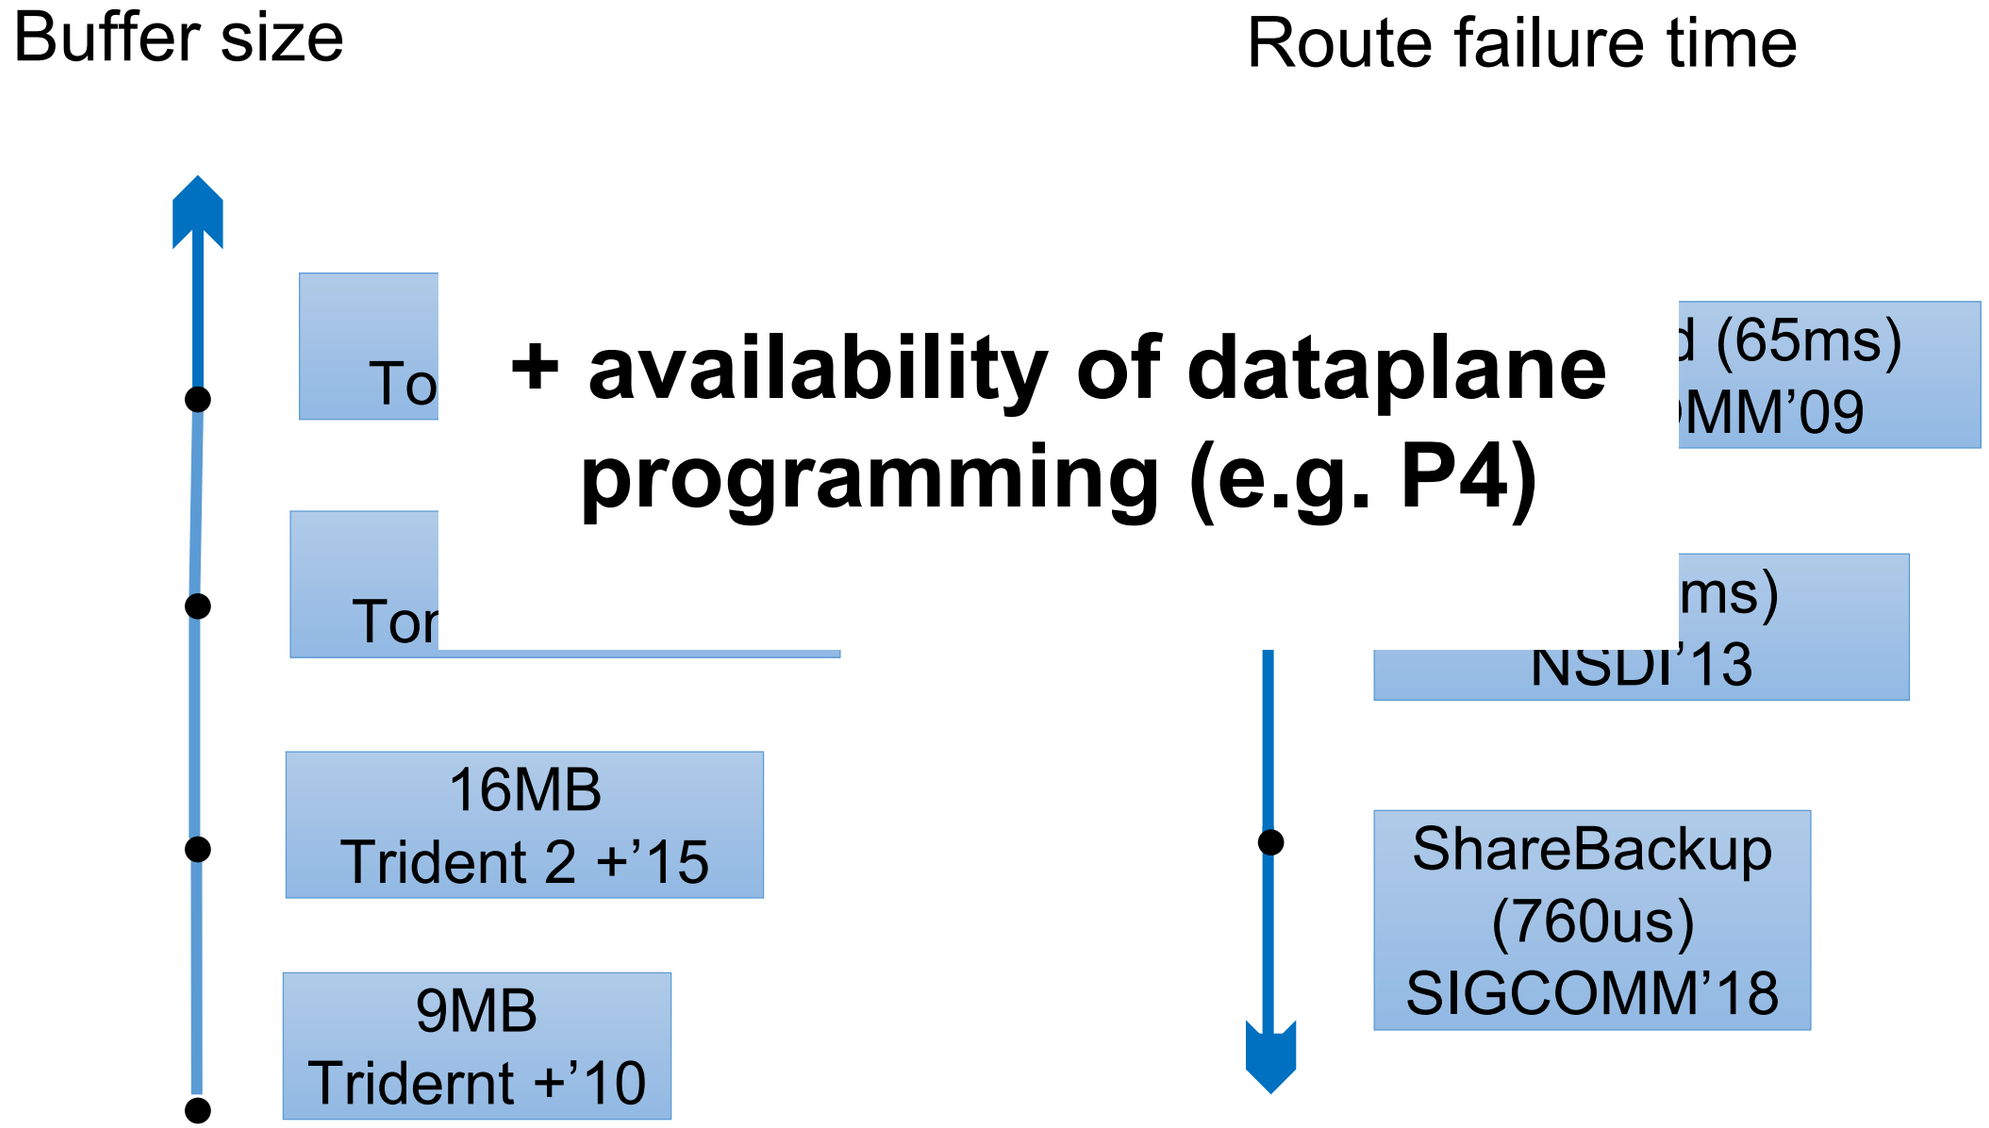

Where and how to cache?

◆Challenges

- In a switch dataplane, the packets can only be stored in the packet buffer with the buffer & queuing engine (BQE).
- The default FIFO queues send out packets as fast as possible.
- No BQE today readily provides the queuing discipline required to realize packets caching with a fixed time.
- BQE does not support custom packet scheduling algorithms.

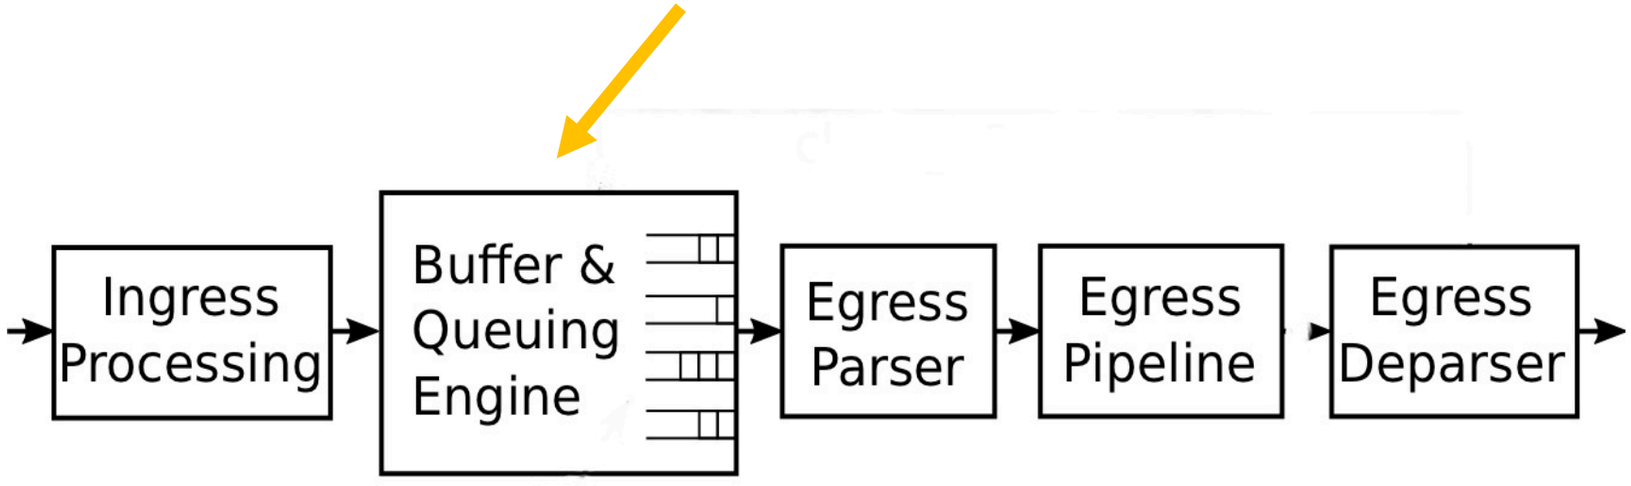

Solution

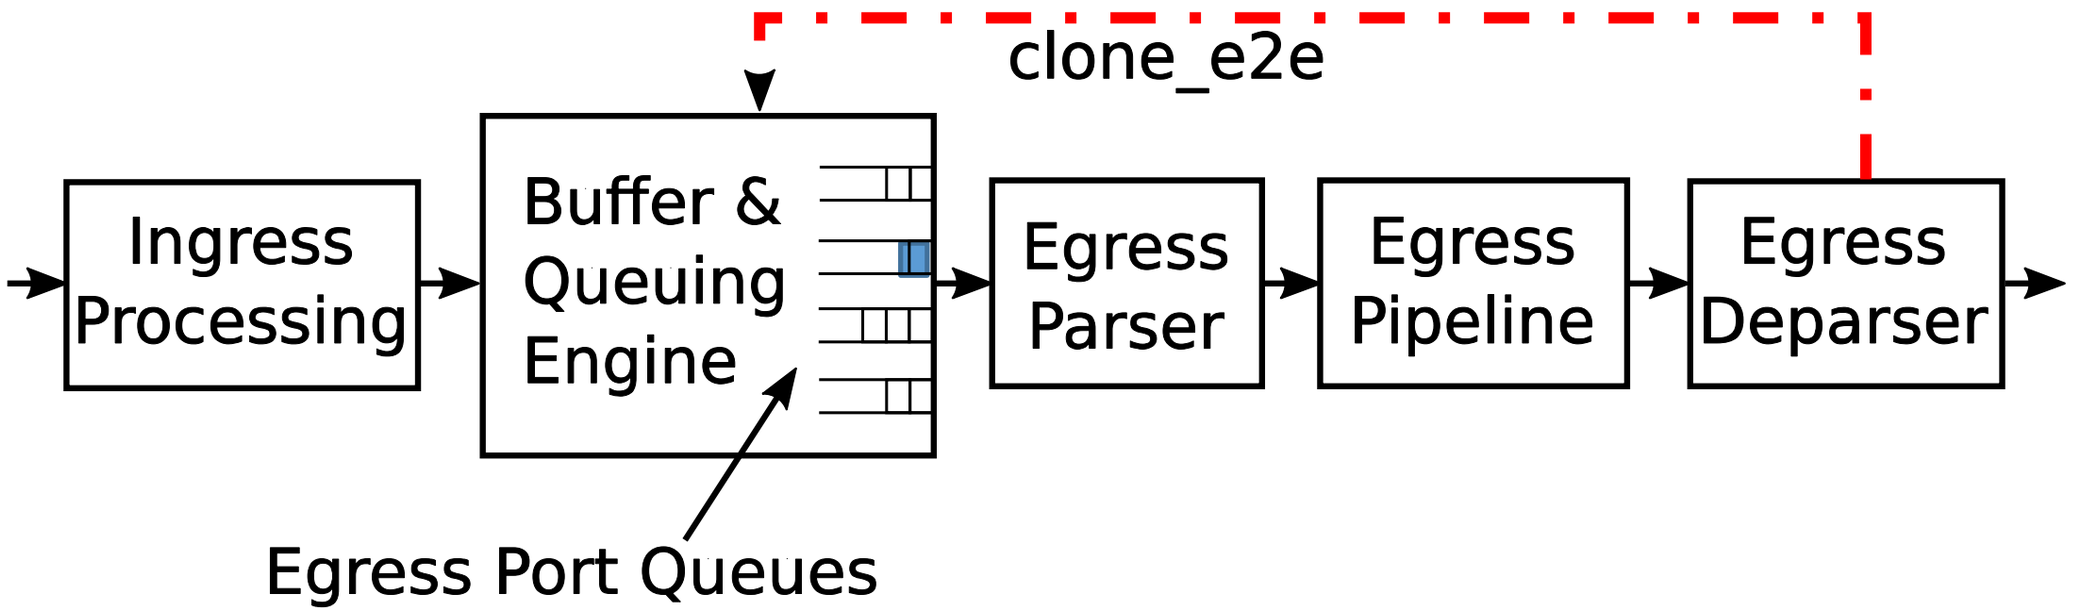

- Keep recent copies of transmitted packets by cloning and then recirculating cloned packets to BQE.
- Supported by the Portable Switch Architecture (PSA)
- Packets are cached for durations sufficiently long to detect link failure and perform route recovery.
- Resend cached packets to new route when it is available.

Challenges

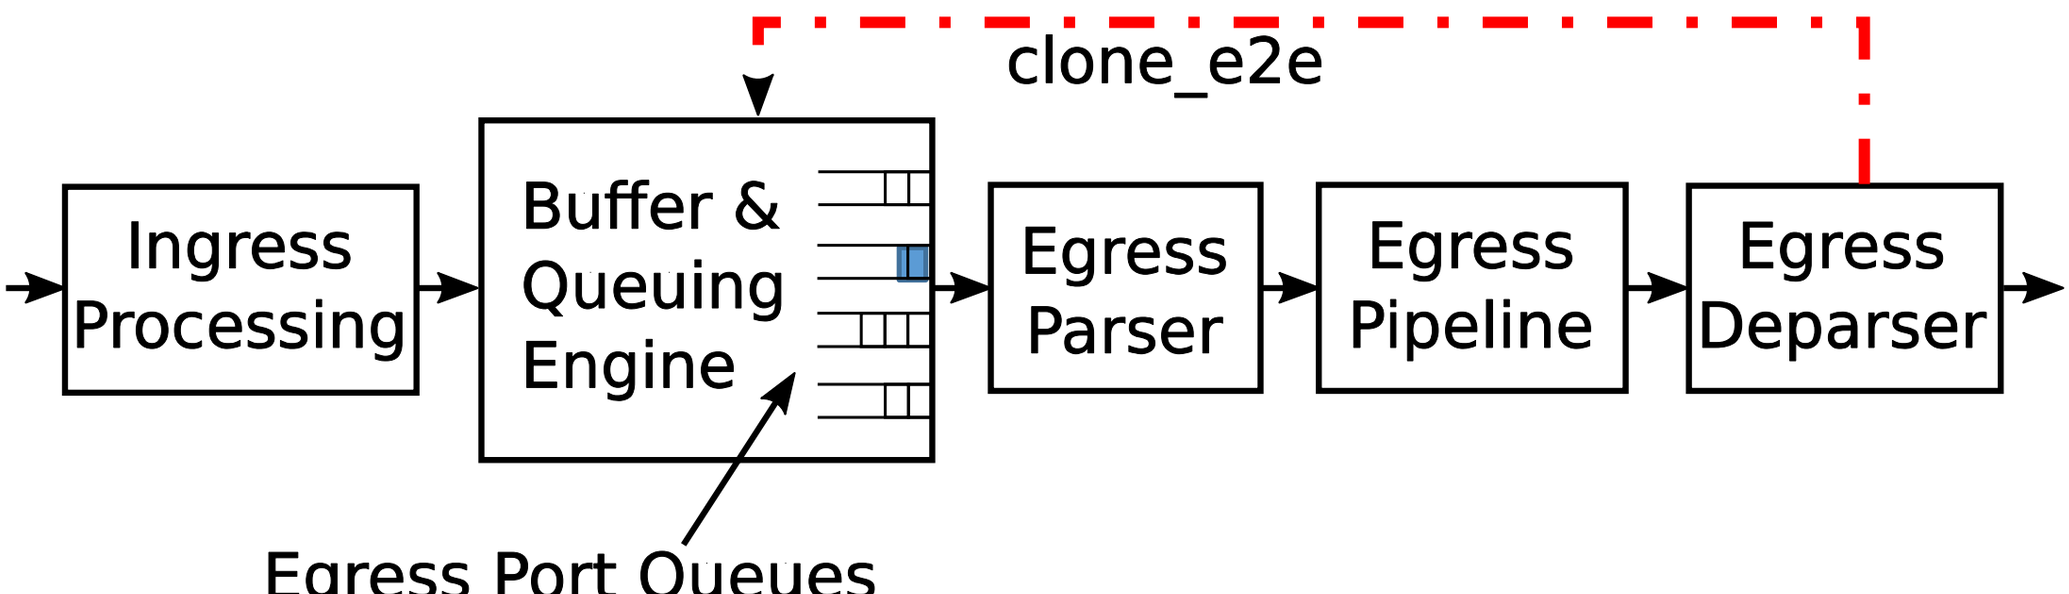

- “Aging” of packets
- Load balancing of circulating packets
- Handle packet reordering

Delay timer

Transmit packet if this is the first/original packet

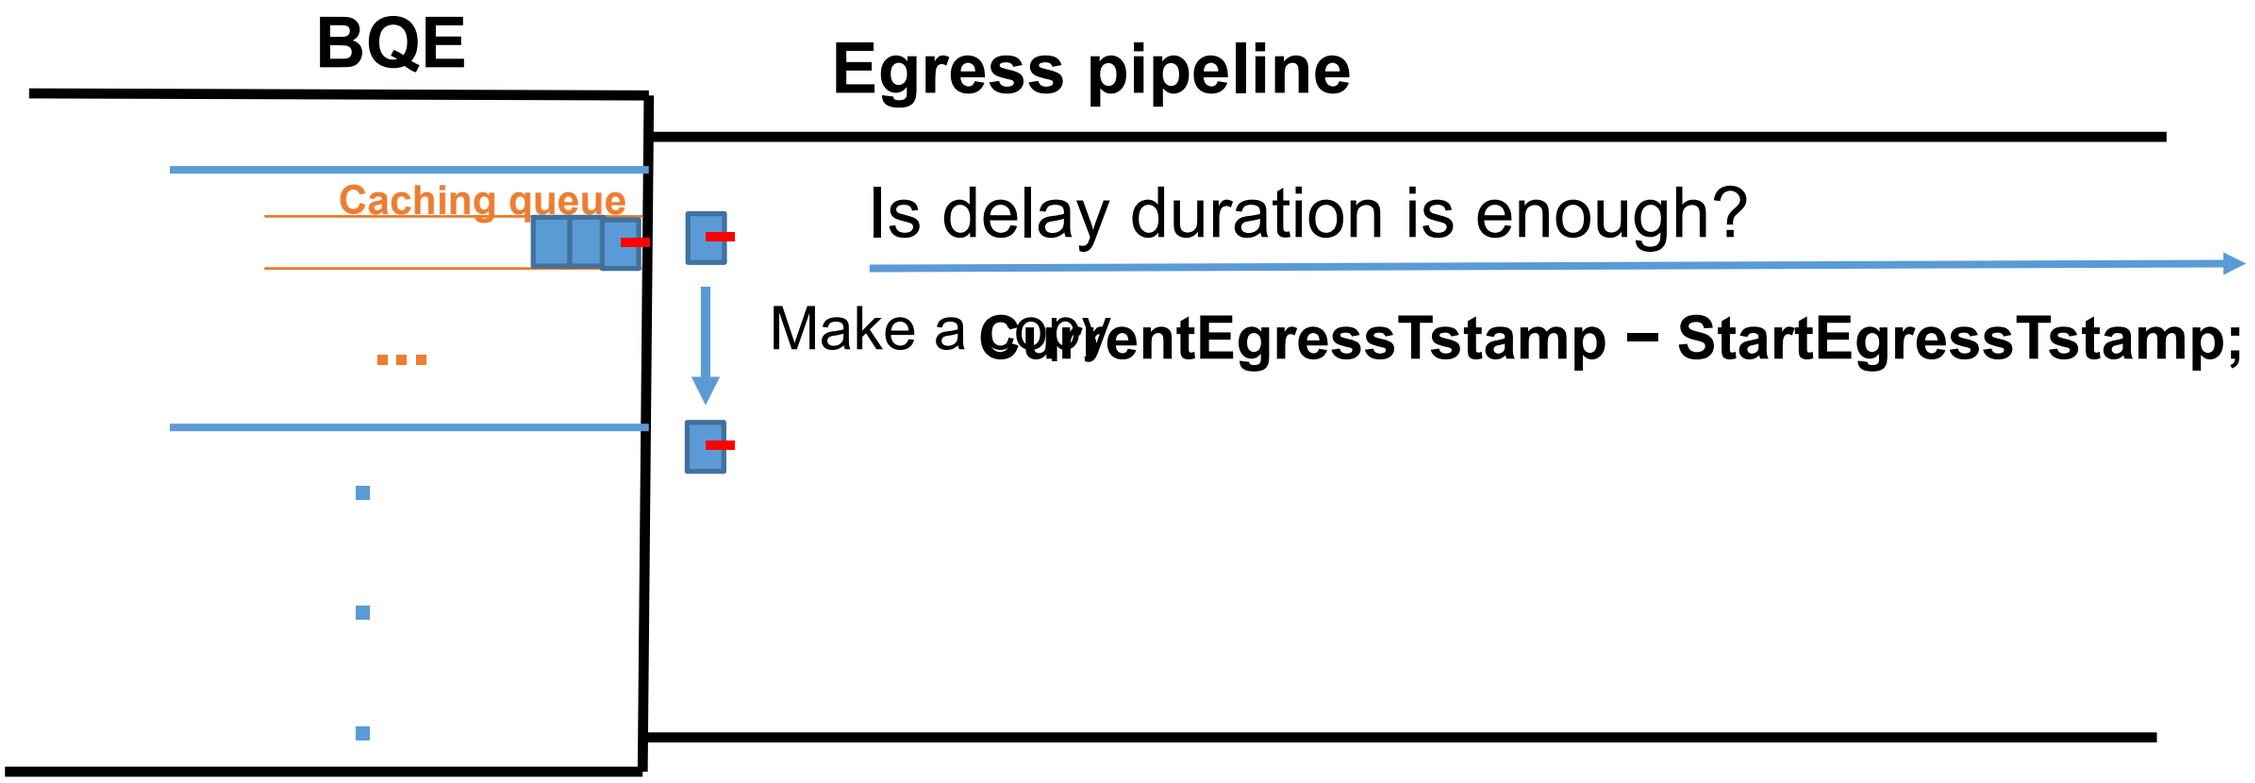

Packet is dropped if it has been cached greater than link detection time

Dynamic queue selection

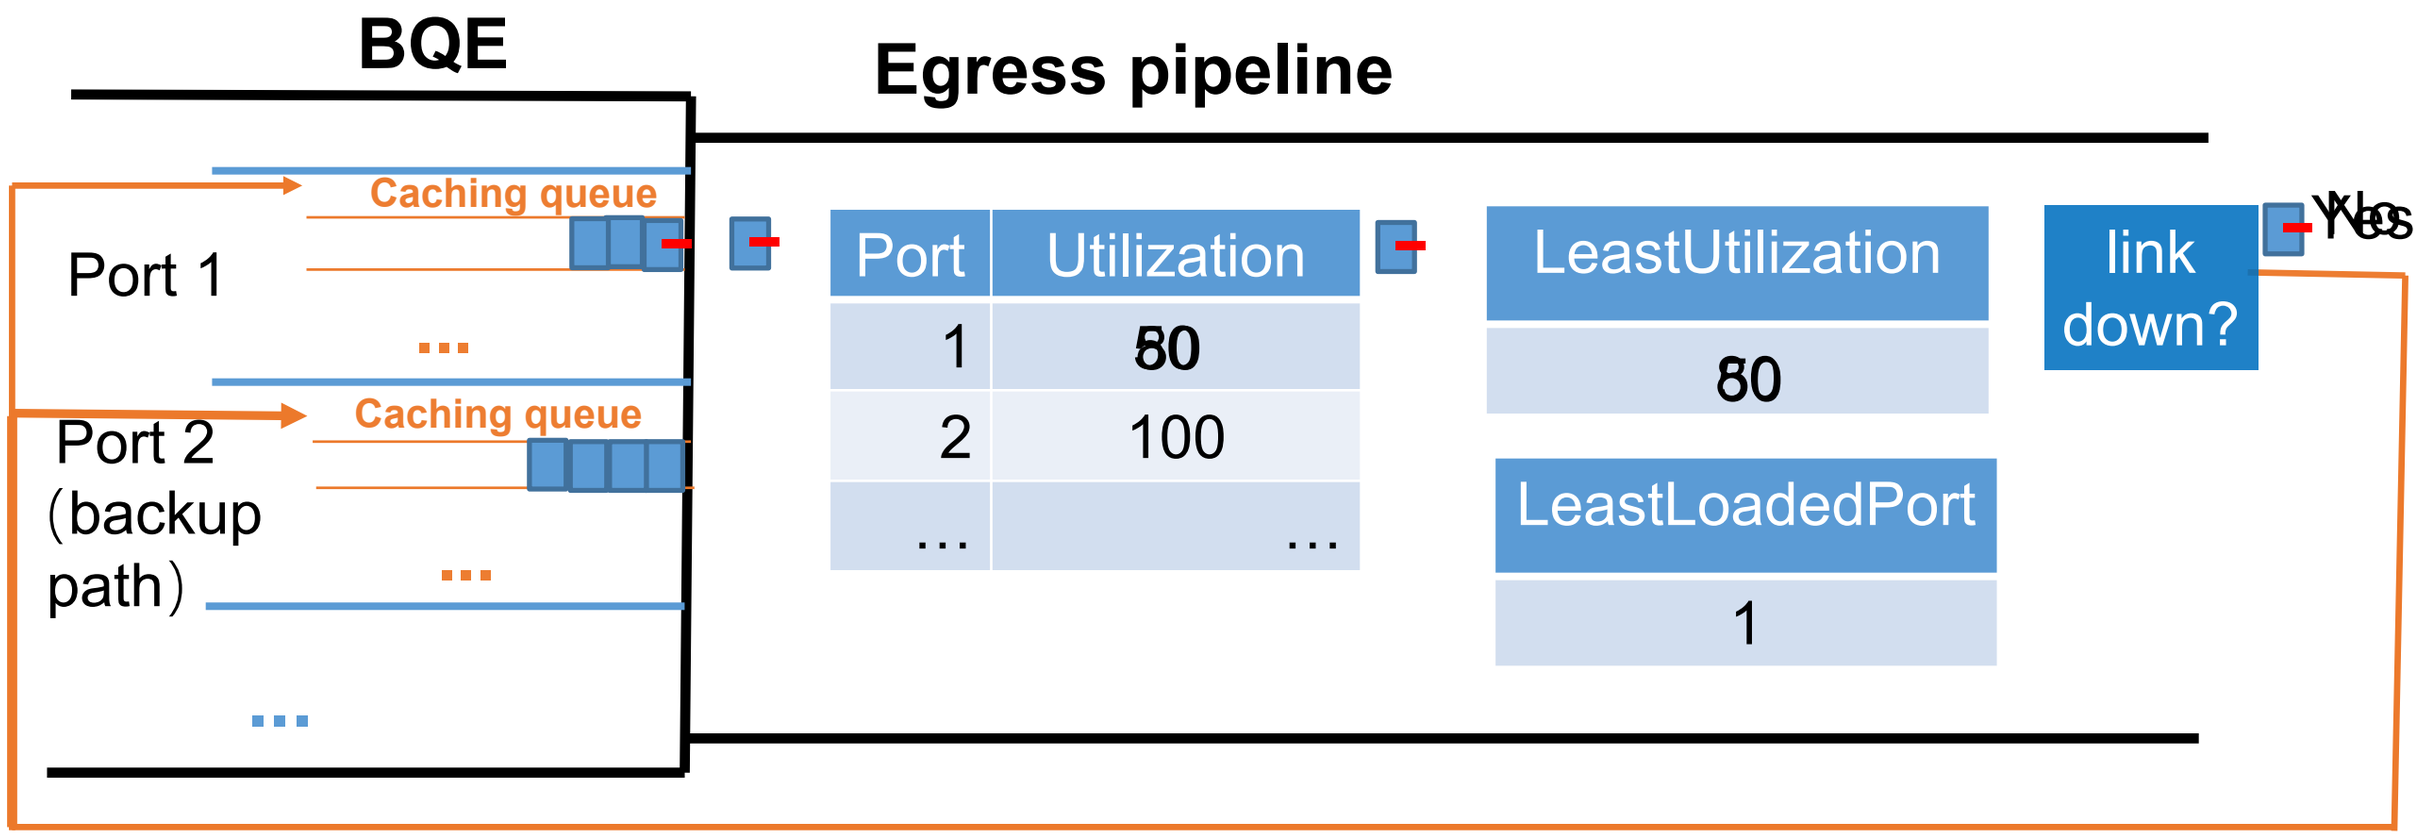

mirroring

Packets from same flow can be cached on different queues

Packet order logic

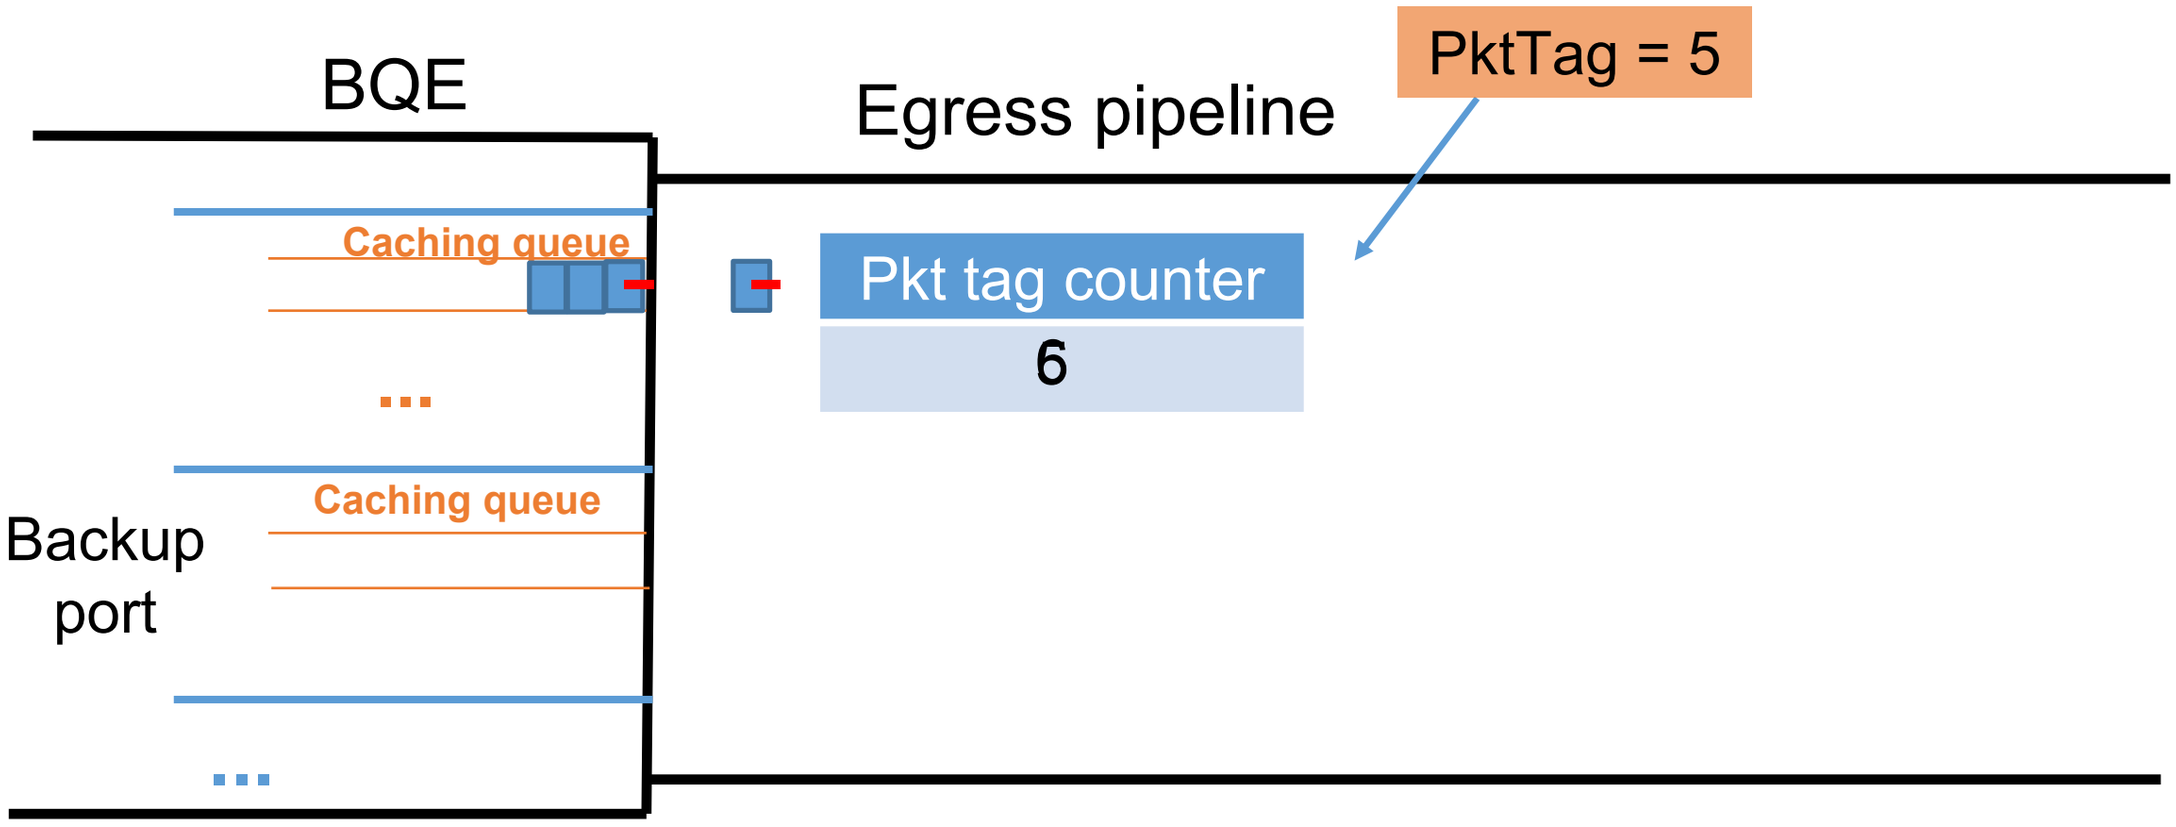

Packet order logic

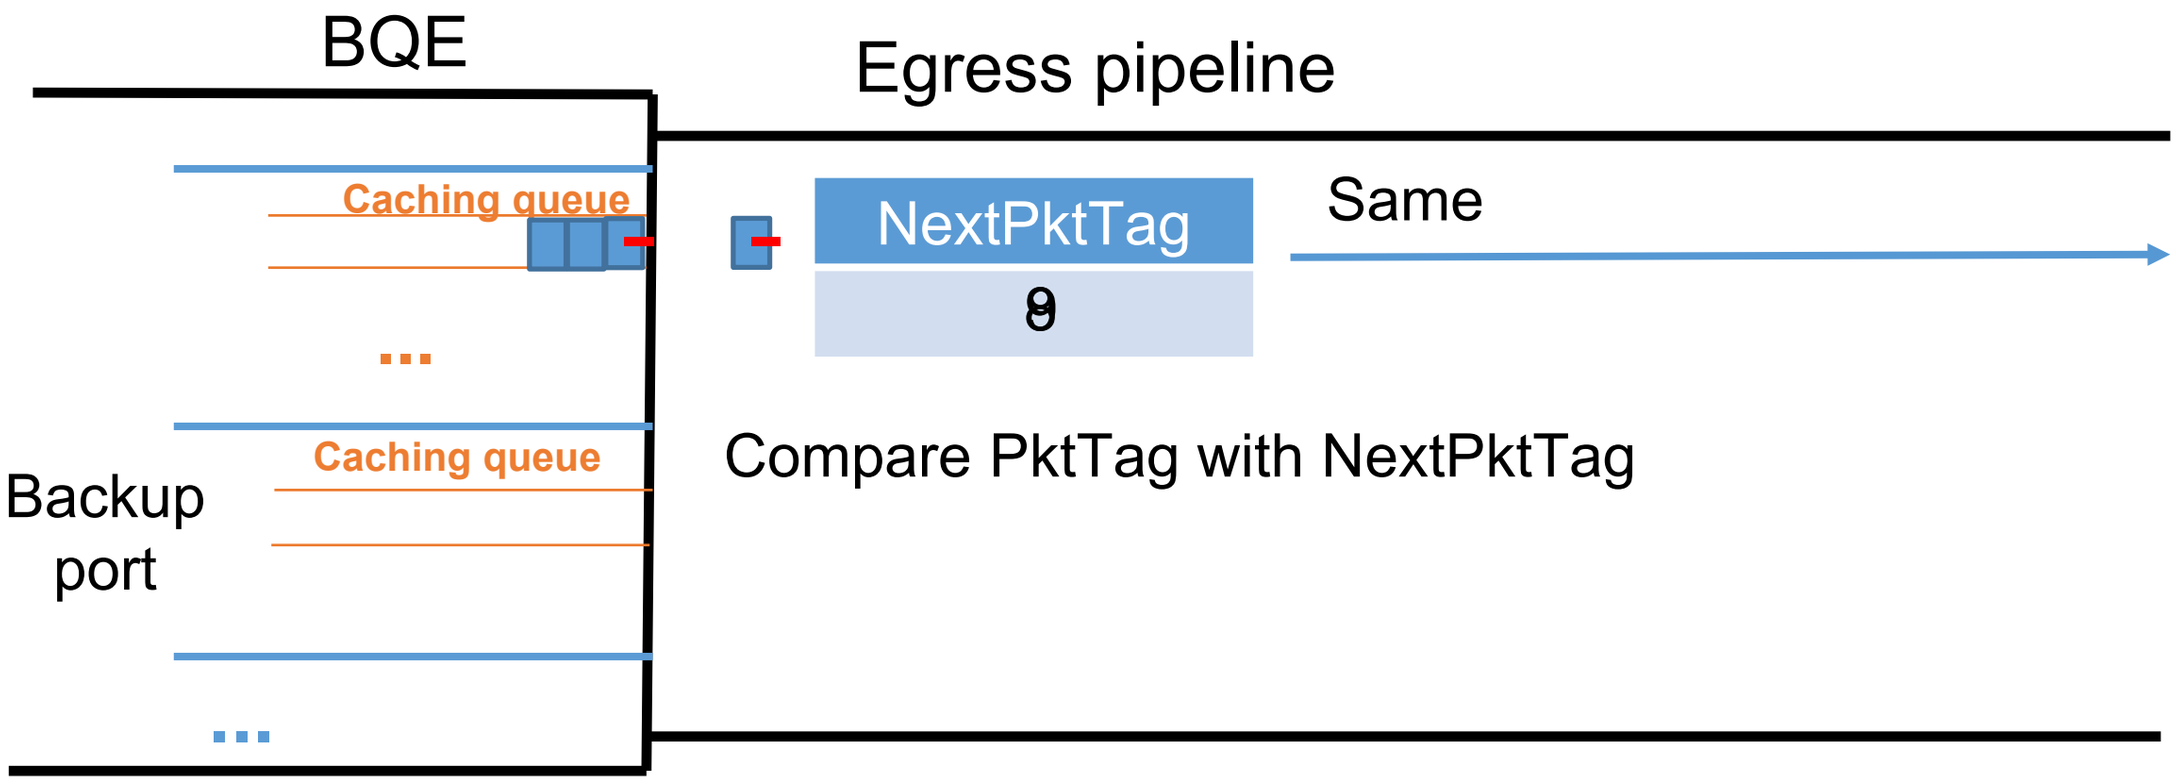

Packet order logic

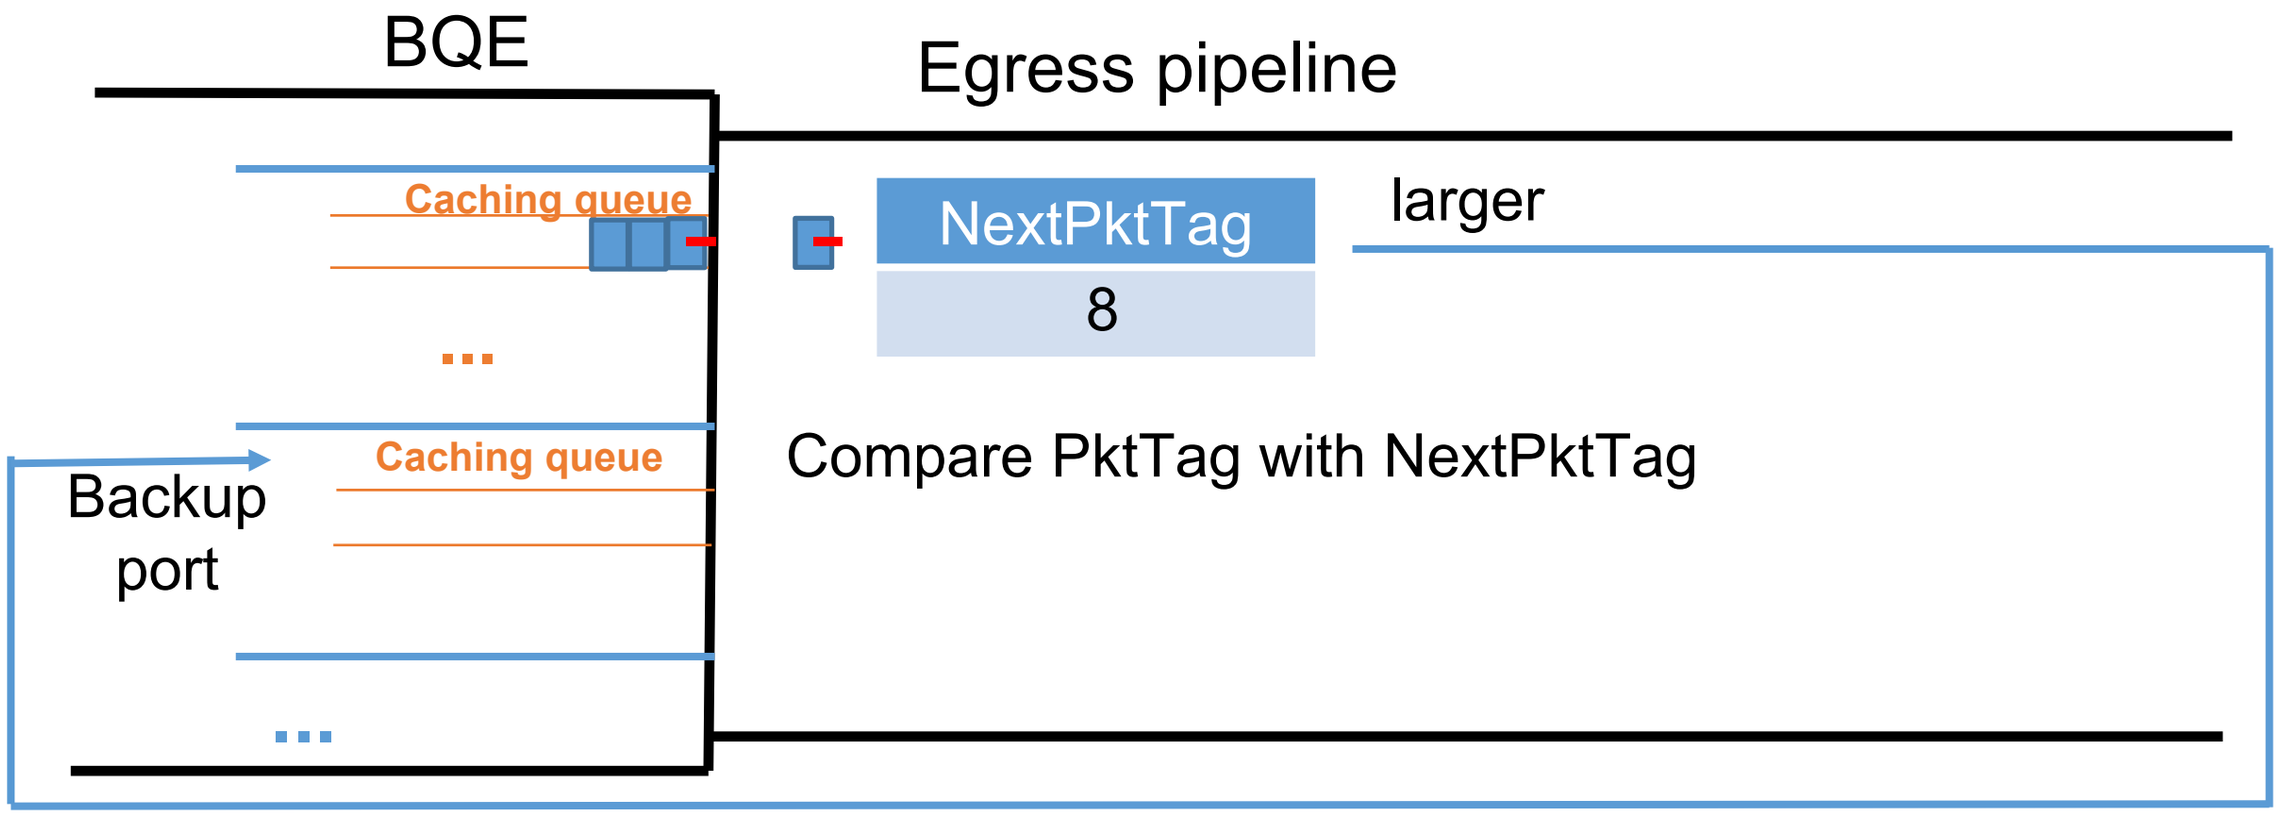

Why it works

- No packet loss
  - ✓ Cache a copy of sent packets for a duration at least equal to the route failure time
  - ✓ Pkt is sent to backup port if new route is ready
- Packets in order
  - ✓ Recover lost pkts based on pkt tag
- Minimize egress processing delays on other flows going through the switch
  - ✓ Select caching queue from multiple ports
  - ✓ Dynamic least loaded port selection
- Complements existing methods of link failure detection and route reconfiguration

Evaluation

- Hardware Testbed
  - Barefoot Tofino switch
  - Intel Xeon servers equipped with Intel X710 NICs
- Trace
  - Web search
  - Data mining
- Schemes compared (SQR implemented in P4)
  - SB’ (simple ShareBackup, 760us route failure time)
  - SB’ + SQR
  - LRR (30us route failure time)
  - LRR + SQR

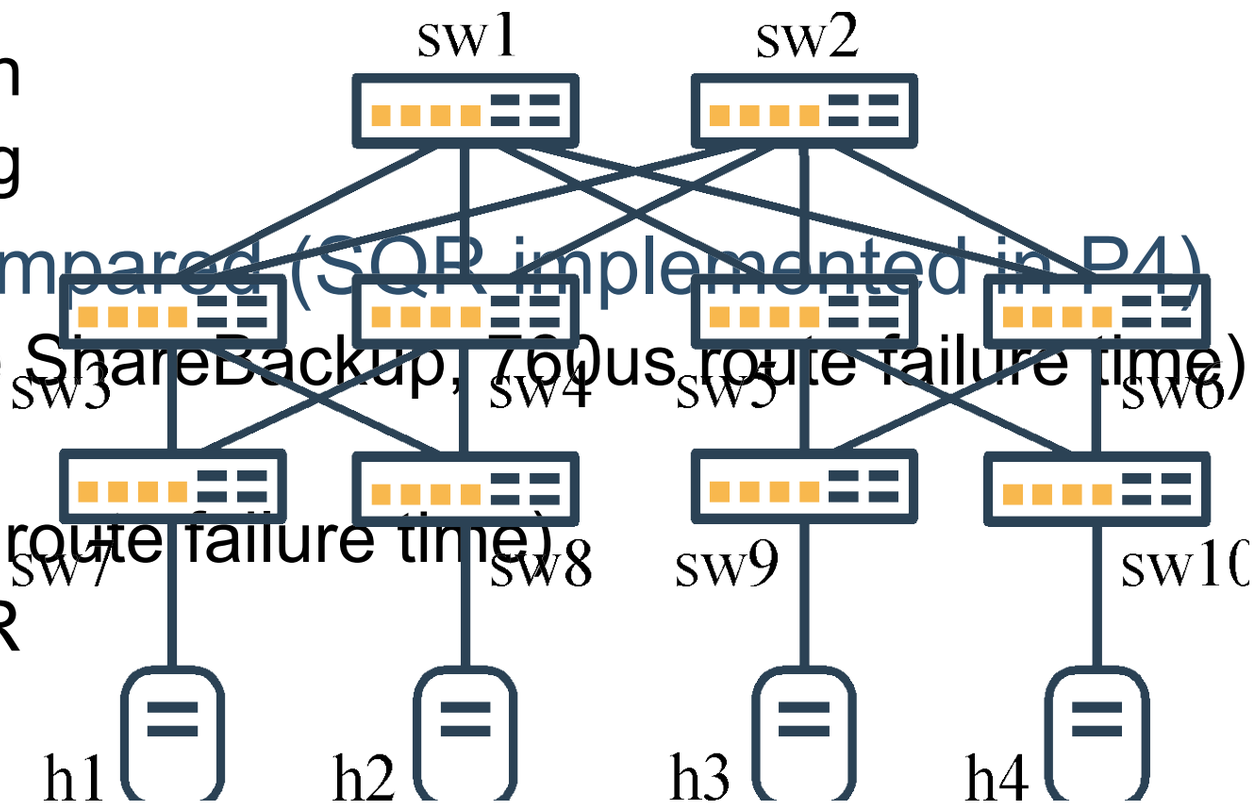

SQR masks link failures from end-point transport

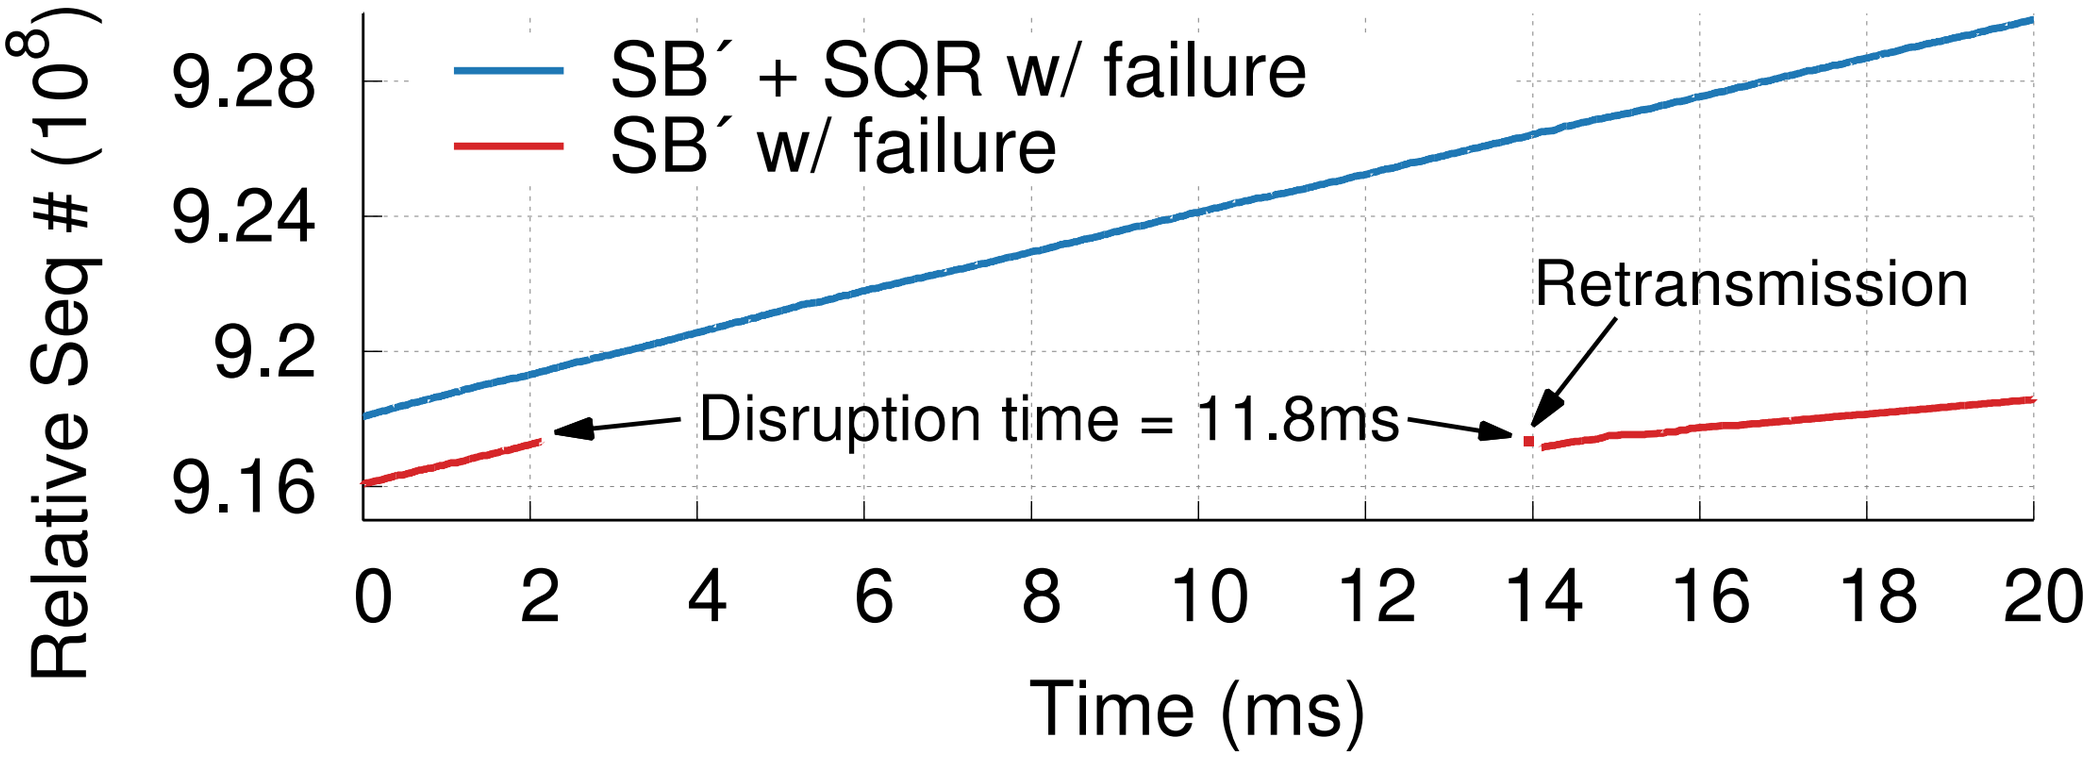

SQR achieves low FCT under link failure

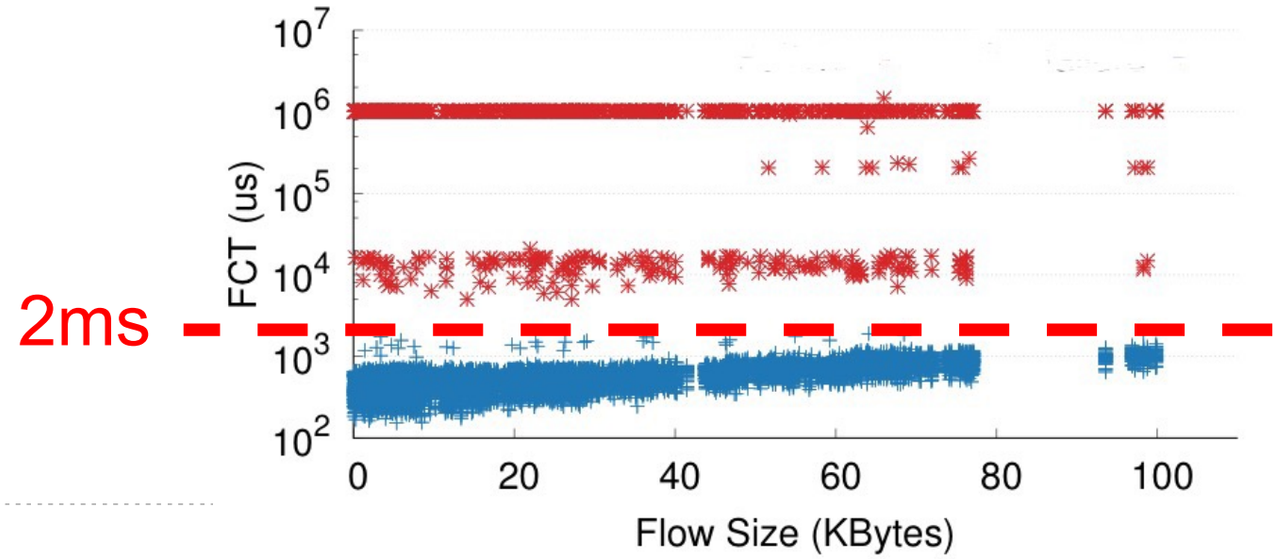

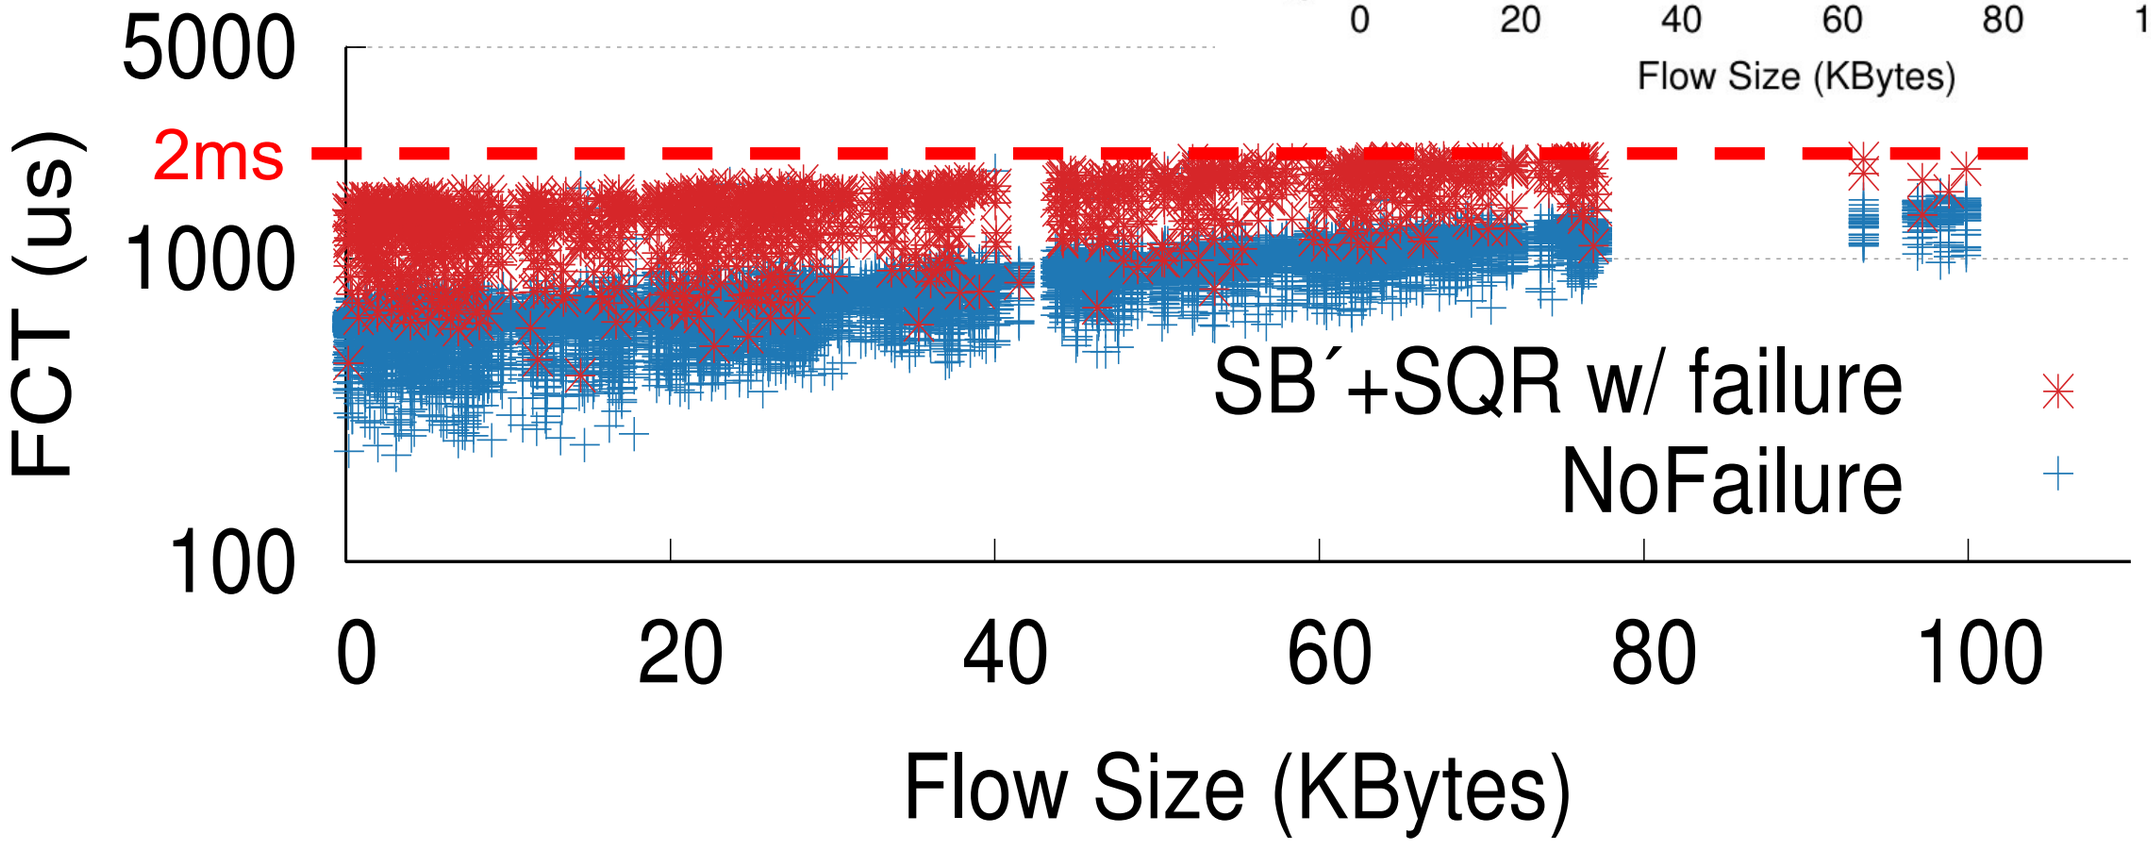

Overhead: Buffer size

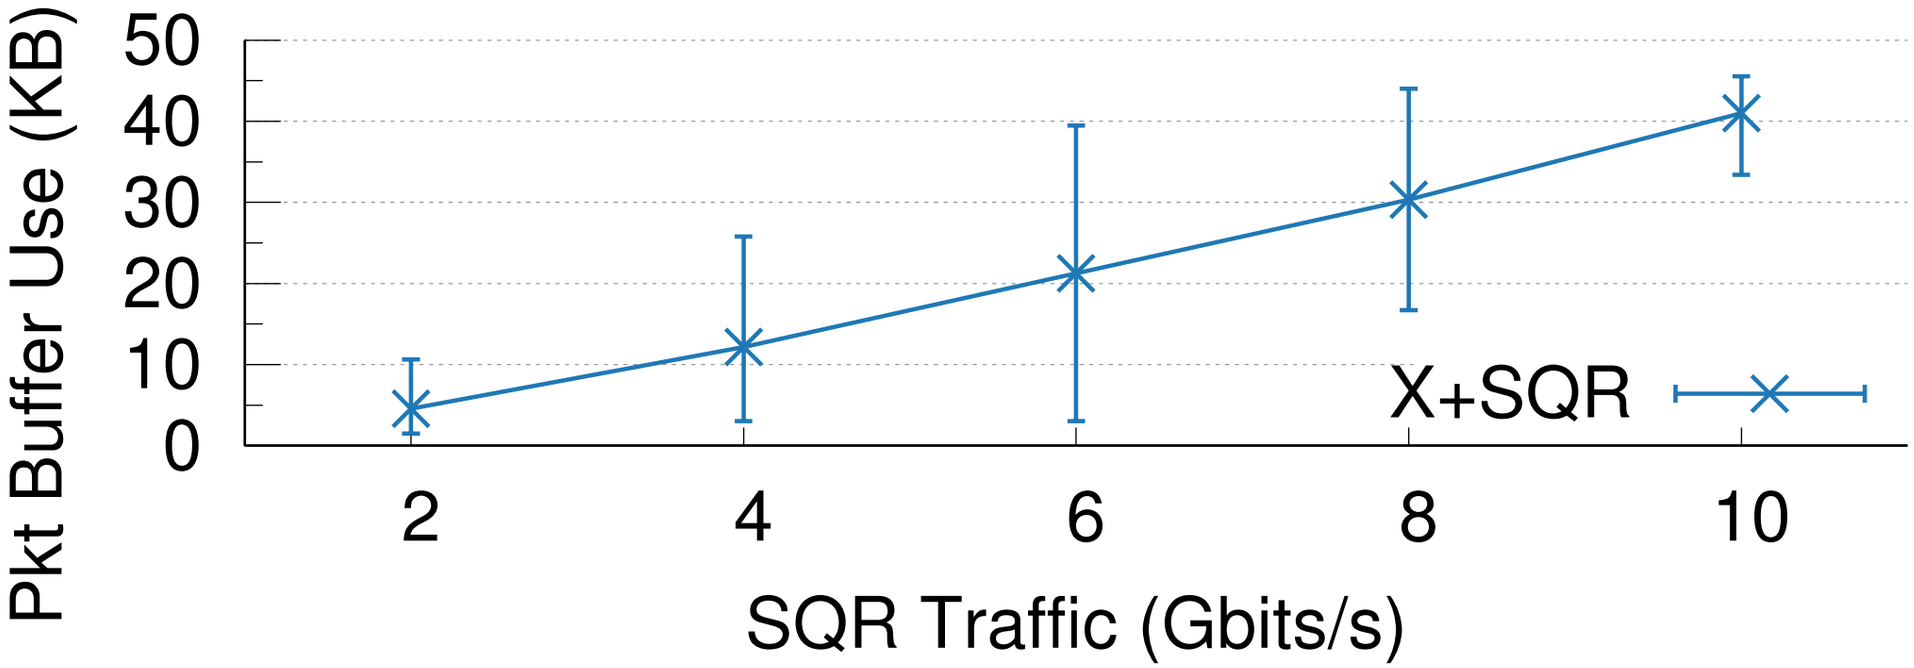

Steady-state packet buffer consumption with 30us link failure detection time

Conclusion

- Design SQR an In-Network packet loss recovery method which keeps FCT low for latency-sensitive flows when there is link failure.

- Eliminate packet loss during link failures and enables handing-off flows seamlessly to alternative paths.

- SQR can be implemented on any programmable ASIC based on Portable Switch Architecture (PSA)


Thank
You
Asante
Maake
Kiitos
Vinaka
köszönöm
спасибо
Dankscheen
Shukria
Dhanyavadagalu
Manana
Dankon
Kam Sah Hammida
ارکش
감사합니다
Dank Je
Mauruuru
Biyan
Matondo
Terima Kasih
Taiku
Tack
Blagodaram
Ngiyabonga
Dziekuje
Chokrane
Diolch i Chi
Grazie
Juspaxar
Arigato
Mochchakkeram
Gracias
Tingki
நன்றி
धन्यवाद
cảm ơn bạn
Gratias Tibi
Ua Tsaug Rau Koj
Bedankt
Dakujem
Grazas
Paldies
Obrigado
Kop Khun Kha
Kia Ora
Děkuji
Nirringrazzjak
ありがとう
Suksama
Hvala
Djiere Dieuf
Rahmat
Nuwun
Matur
谢
Welalin
Di Ou Mèsi
Eskerrik Asko
Misaotra
Danke
Merci
Go Raibh Maith Agat
Salamat
ขอขอบคุณคุณ
Najis Tuke

Impact of SQR Traffic

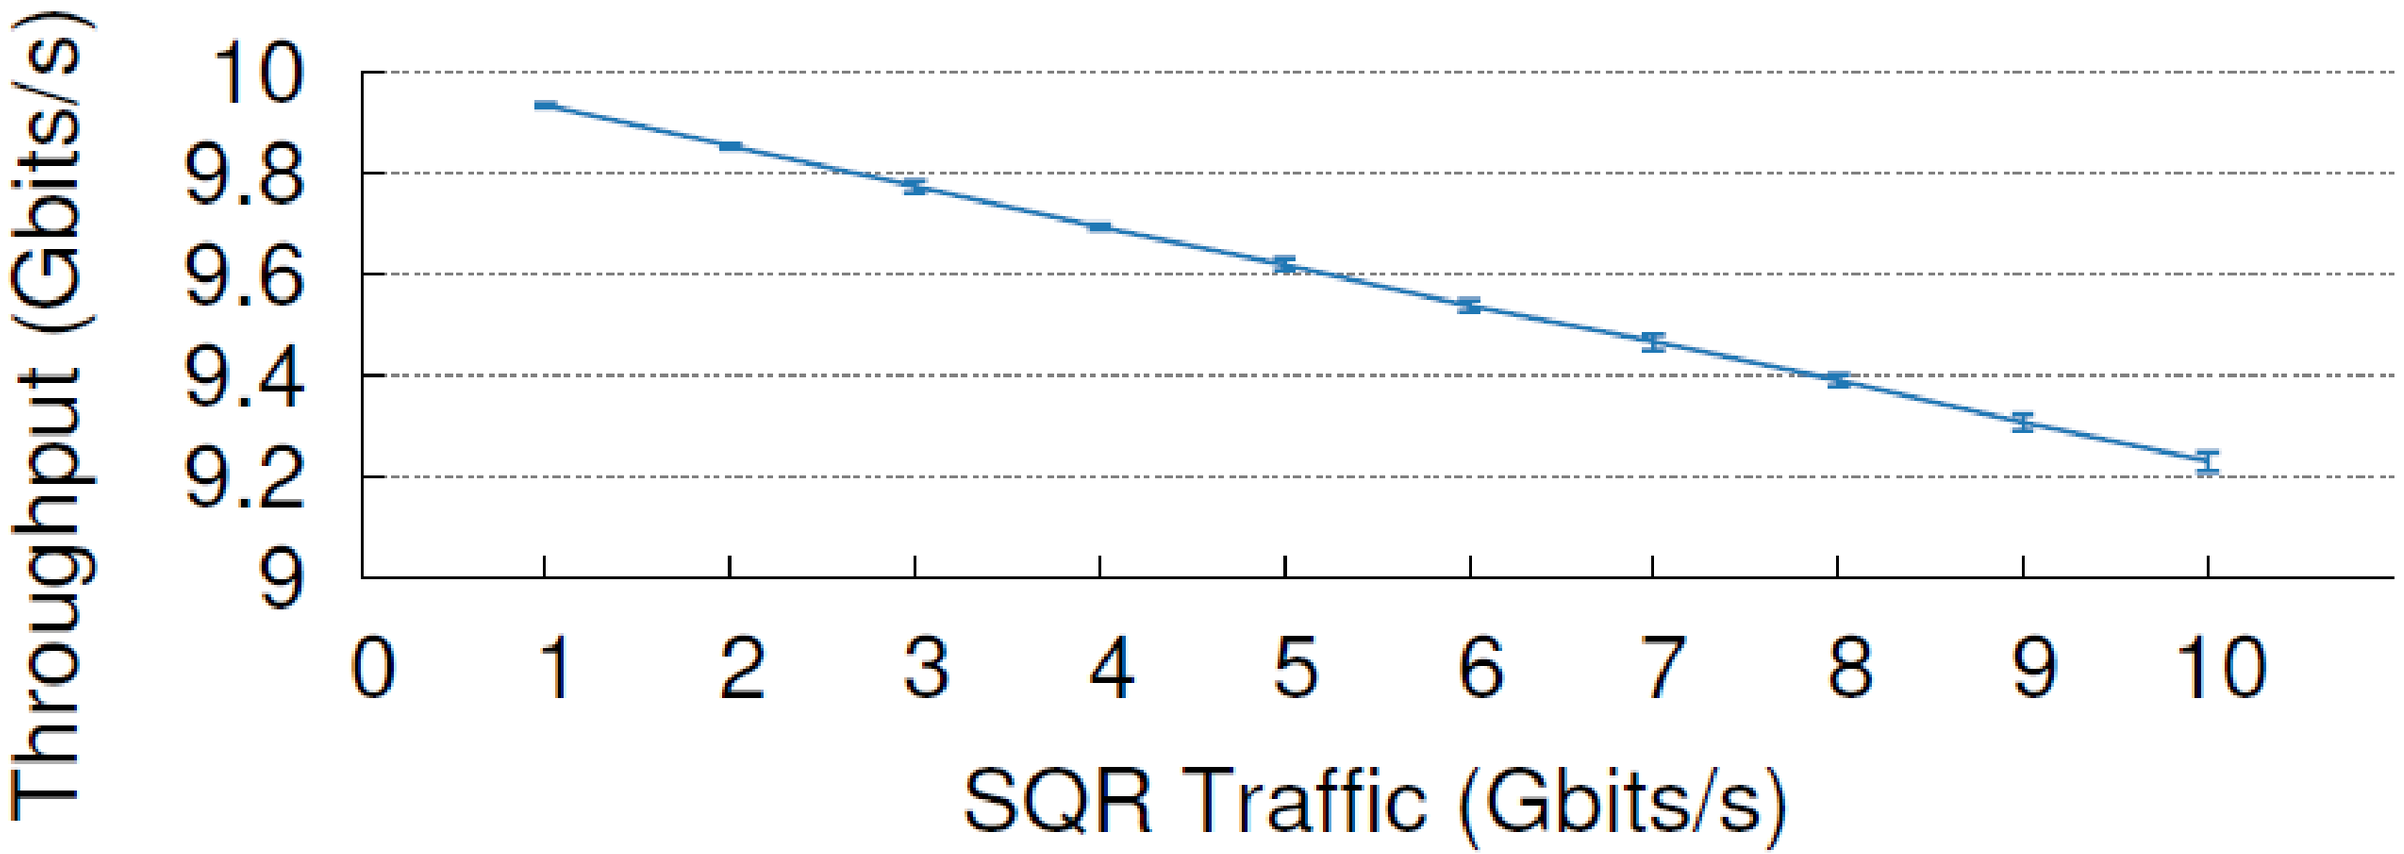

Overhead: Egress processing