TRANSPACKET

Zero packet jitter aggregation and priority mechanisms

09.03.2015

Steinar Bjørnstad

Steinar.bjornstad@transpacket.com

Outline

- Low packet-jitter Packet Delay Variation (PDV) applications and requirements
- Mechanisms for reducing latency and packet-jitter
  - FUSION time-window
  - FUSION guaranteed service stream aggregation
- Experimental results using FUSION on GE and 10GE
  - Adding lower quality traffic to 10GE guaranteed stream
  - GE aggregation into 10GE

FUSION history

- Year 2000: Optical networks becomes packet switched
  - Moving from Circuit to packet switching
  - Maintaining properties from circuit switching is desirable
- Carriers wondering how to maintain circuit properties when moving to IP and Ethernet based networks

Integrating packet and circuit

- Integrated hybrid optical network (IHON)
  - Properties from Circuit and Packet networks combined
  - Still pure packet based, no TDM
- FUSION networks
  - IHON known from academic literature: Published in major IEEE conferences and journals
  - Not standardized
  - Commercialized as FUSION networks (TransPacket)

Carrier pain

- Wavelength and OTN services have a high production cost
  - Occupies a physical wavelength or TDM channel resource in the network
  - Wavelengths/TDM channels are limited resources
  - Wavelengths/TDM channels occupies resources end-to-end – no intermediate additional aggregation possible
- Ethernet or VPN service preferred as compromise
  - Lower production cost (oversubscription/statistical multiplexing)
  - Does not offer transparency and performance (especially latency) as for wavelengths and OTN channels

Optical networks

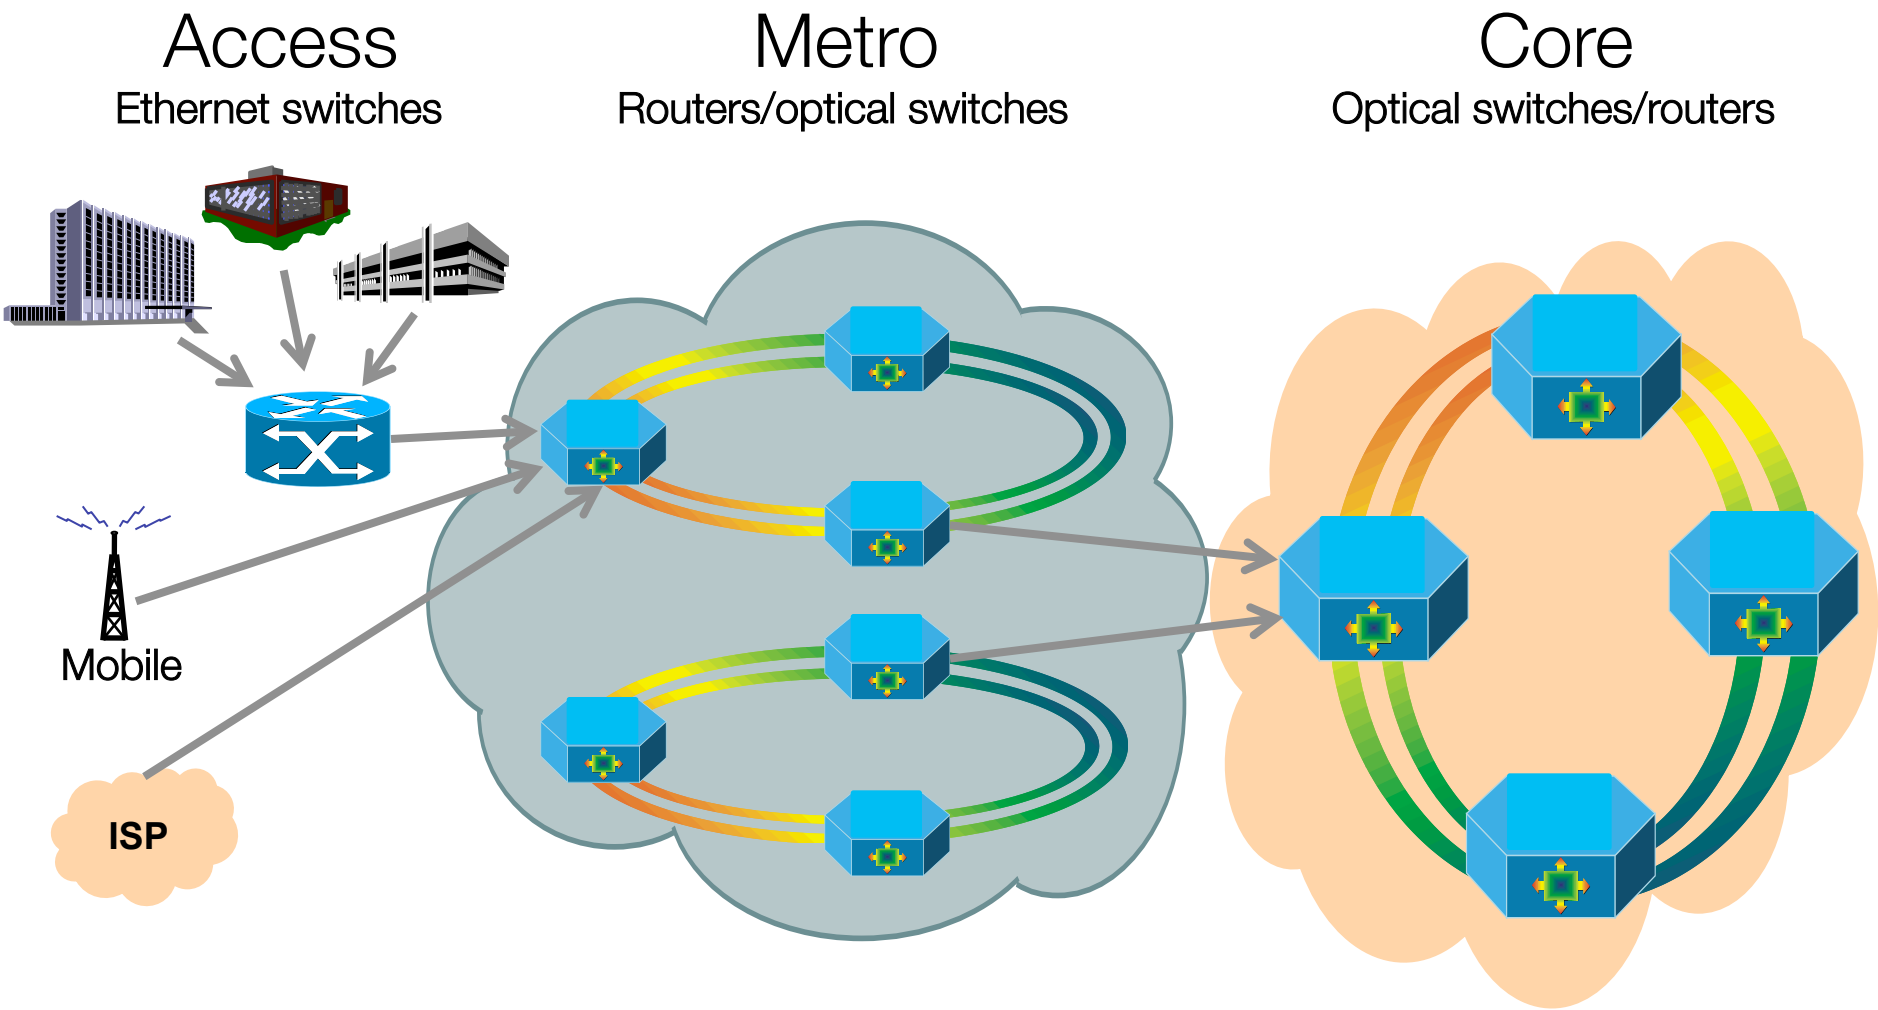

Wavelength services are costly

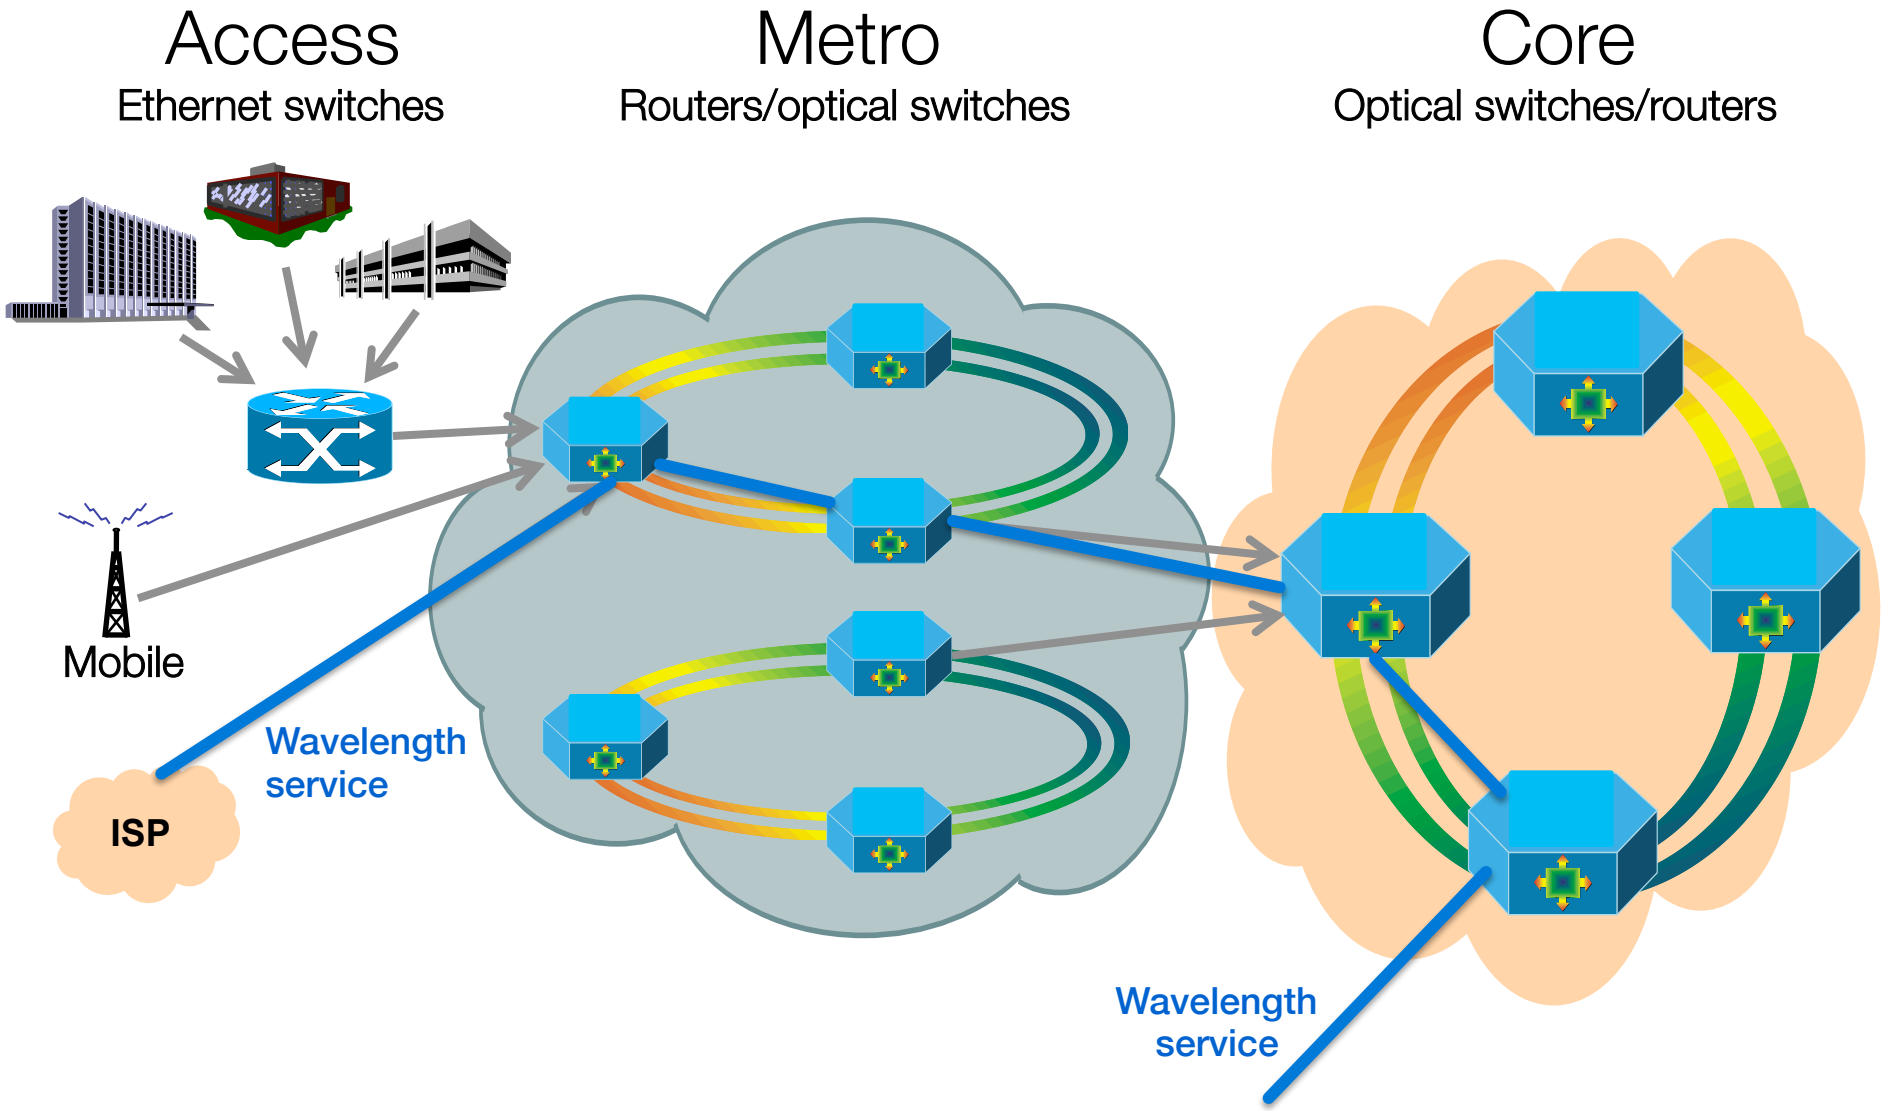

Ethernet Virtual wavelength services: Sub-wavelength granularity

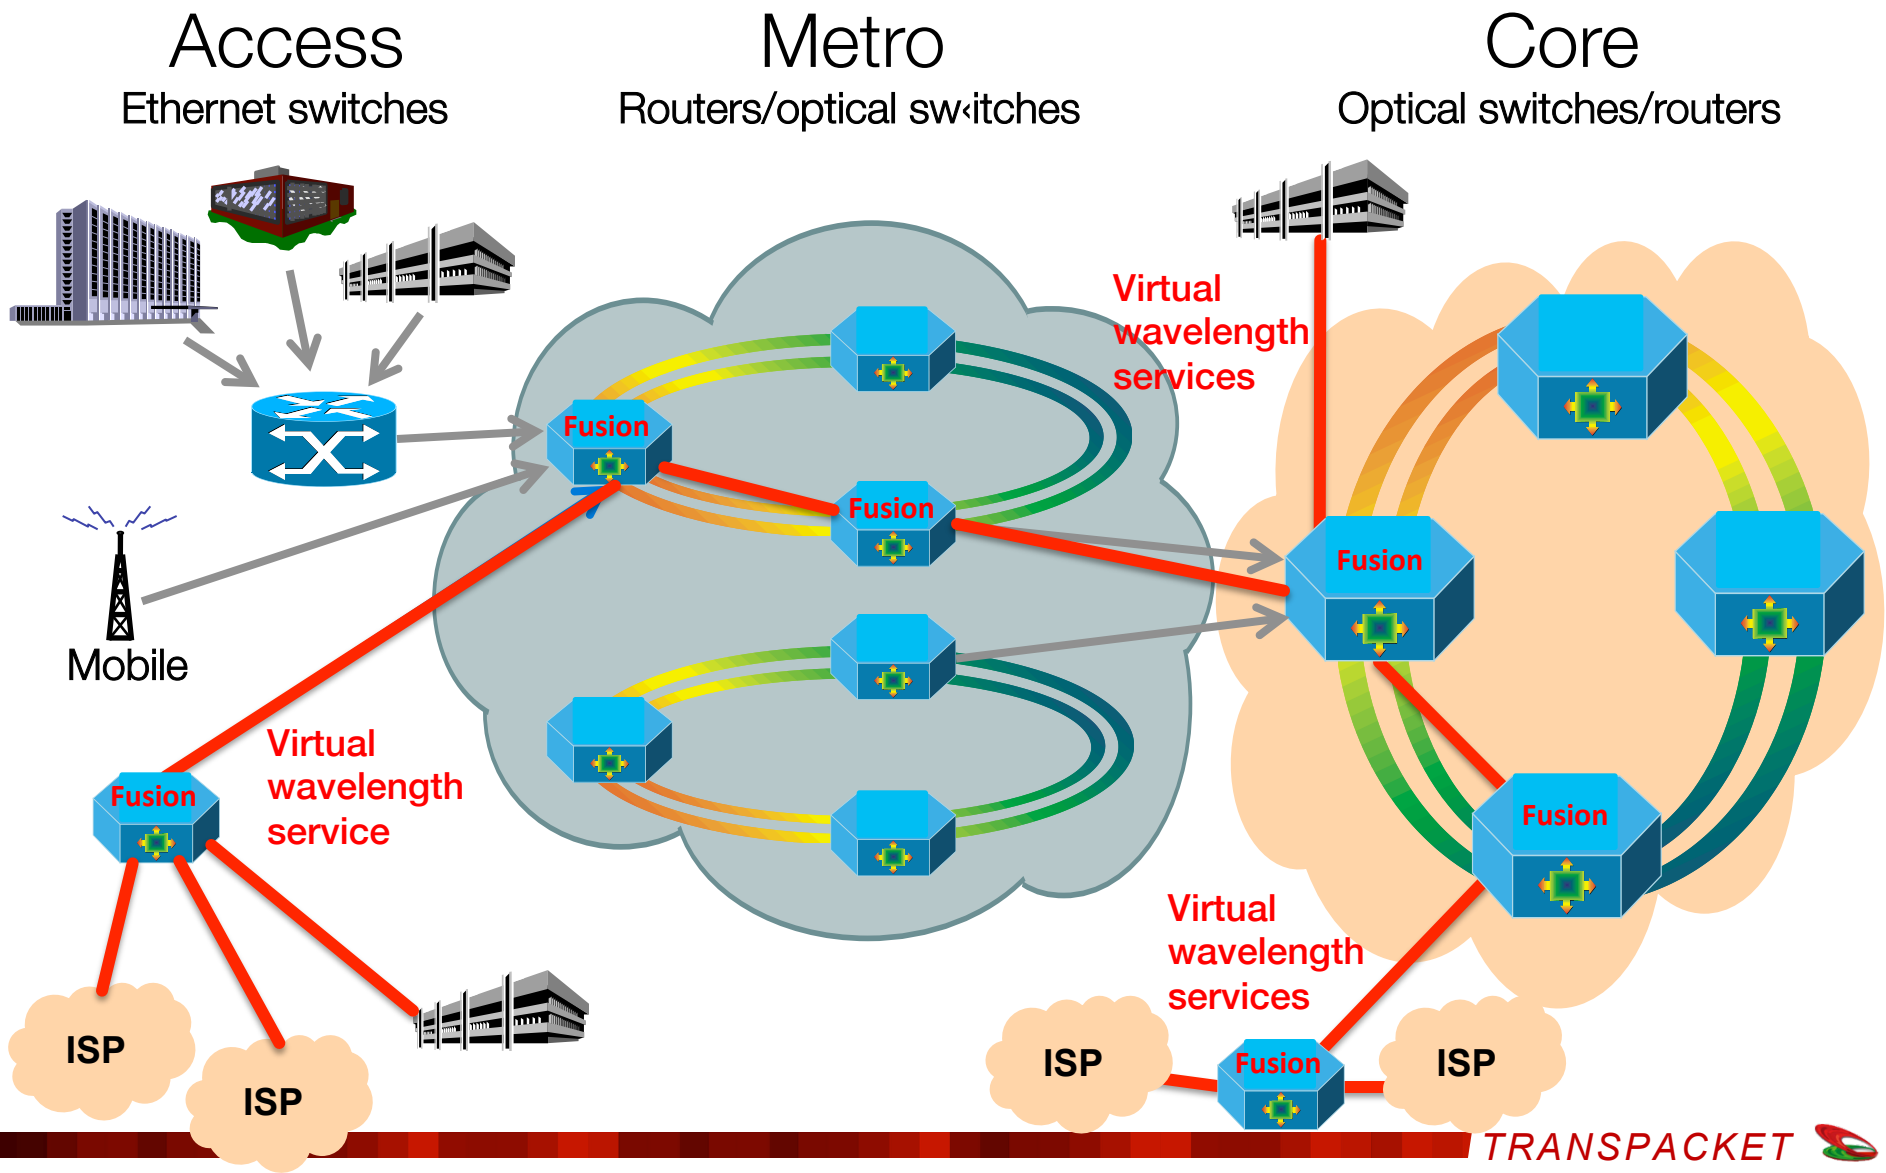

FUSION applications and requirements

- Ethernet Virtual wavelength services
  - Performance close to circuit switching (OTN)
  - Ethernet throughput performance
  - No packet loss due to contention
  - Fixed latency (Packet-jitter in the ns range)
  - Low latency, but **fibre is 5 microsecond/kilometer**
- Mobile backhaul and fronthaul
  - Transparent transport of IEEE1588 packets in mobile backhaul: Strict demands on PDV
  - Radio over Ethernet in mobile fronthaul: Extreme demands on packet-jitter and strict on latency

Packet-jitter and latency demands

- Mobile backhaul
  - Using IEEE1588 for synchronization
  - Average jitter important to keep low
  - Sync. Accuracy in the microsecond range
- Mobile fronthaul - CPRI over Ethernet
  - Peter Ashwood Smith, IEEE 802, July 2014:

REQUIRED

- 100us – Maximum one way Delay between Antenna and Compute
- 65ns – Maximum variation in Delay (Jitter).
- 1-10G – Throughput per antenna (compression possible).
- 10-12 – Maximum Bit Error Rate

OBSERVED

Average latency = 3us per hop

Peak Jitter = +/- 2500ns per hop

Packet-jitter and latency demands

- Mobile backhaul
  - Using IEEE1588 for synchronization
  - Average jitter important to keep low
  - Sync. Accuracy in the microsecond range
- Mobile fronthaul - CPRI over Ethernet
  - Peter Ashwood Smith, IEEE 802, July 2014:

REQUIRED

- 100us – Maximum one way Delay between Antenna and Compute
- 65ns – Maximum variation in Delay (Jitter).
- 1-10G – [illegible]
(compress[illegible]
- 10-12 – [illegible]

OBSERVED

Average latency = 3us per hop

Peak Jitter = +/- 2500ns per hop

Peak-Jitter is the main challenge!

Mechanisms for reducing latency and packet jitter (PDV)

Preemption minimizes delay

FUSION minimizes packet-jitter

(patent/application: PCT/GB12/50202)

Without Hold and Release

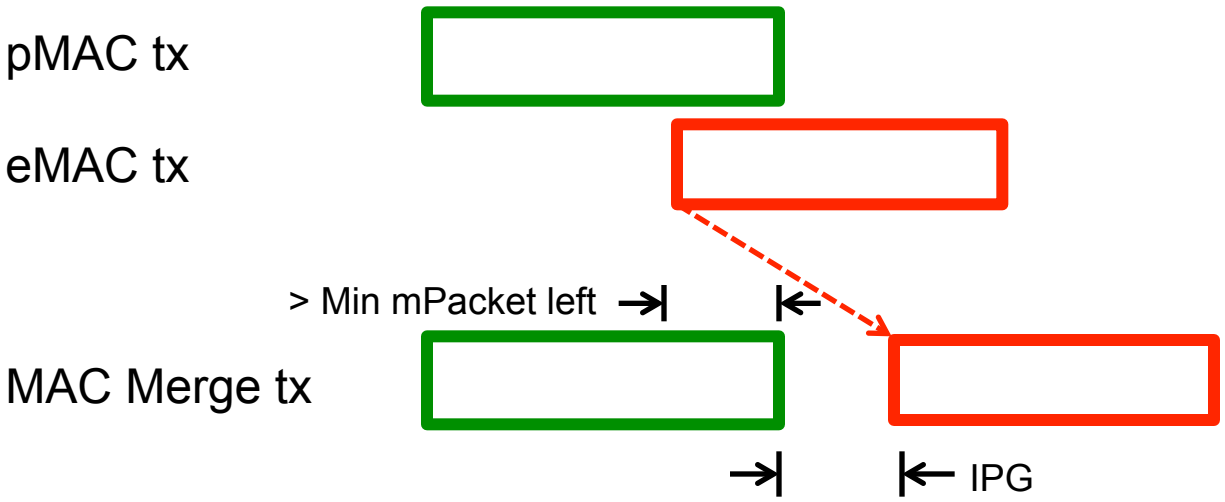

- Preemption isn't instantaneous.
- Packets with less than min packet size left to transmit or packets less than 123 octets can't be preempted.
- In many use cases, this delay is short enough but not in all cases.

Underutilized circuit (wavelength): Fusion fills it!

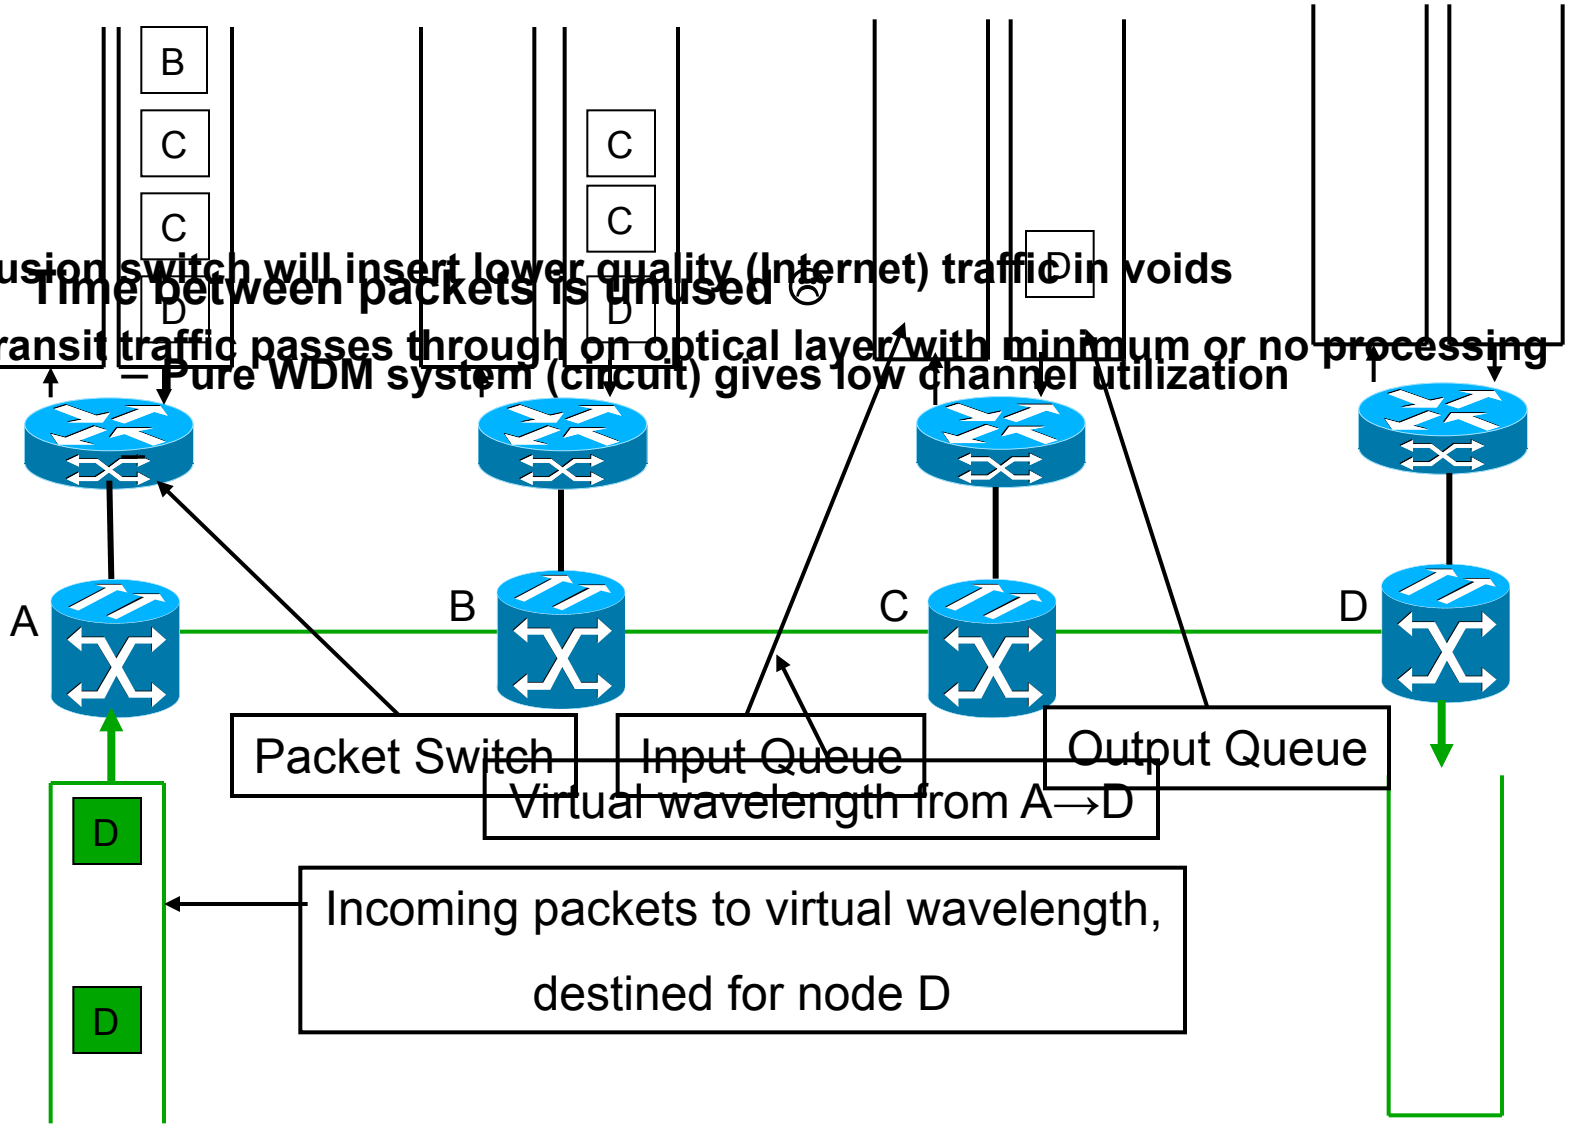

Time-window: Size of packet gaps are known

Lower quality streams

SM input queues

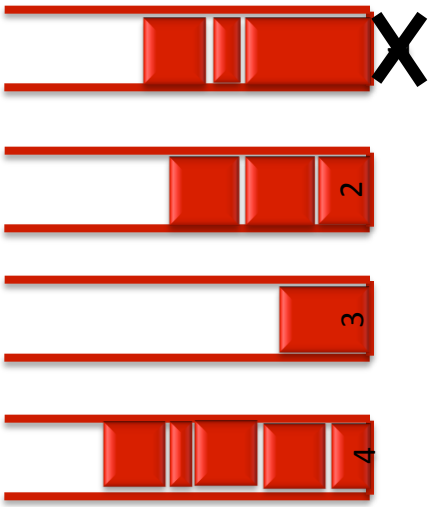

- Round-robin on the SM input queues
- Schedule SM packet that fits the gap
- Reduces head of line blocking!
- One queue per port : Avoids packet re-ordering

GST: Guaranteed Service Transport stream packets

Delay and packet-jitter

- Time-window detection introduce a fixed delay on high-priority traffic, but packet-jitter is zero.
- The introduced delay = Duration of one lower-priority MTU, e.g. 1518 Bytes
- Preemption introduce a delay **and packet-jitter** of 123 Bytes + IFG = 137 Bytes
  - Approx. 1/10 of the FUSION time-window approach

Delay and packet-jitter but how much?

- Assuming MTU = 1518 Bytes
  - 1 Gb/s = 12 microseconds
  - 10 Gb/s = 1.2 microseconds
  - 100 Gb/s = 120 nanoseconds
- Fibre delay is 5 microseconds/km
- Optical networks are typically 10`s of kilometers and minimum 10 Gb/s
- In optical networks one MTU delay is typically low compared to transmission delay.

Pre-emption versus time-window

- FUSION Time-window
  - Zero packet-jitter on high priority traffic
  - No packet fragments – No “Tax”
  - Retransmission mechanisms not required
  - Higher throughput on lower-quality traffic than pre-emption
  - Easy to introduce gradually, node-by-node
- Preemption
  - Lower latency than FUSION Time-Window

FUSION Guaranteed stream aggregation

- Alternative to 802.1Qbv time gated queues
- Principle:
  - Do a photocopy of each of the incoming lower bitrate streams
  - Place the copy into dedicated (virtual) time-slots, i.e. a container in a higher bitrate stream
  - The copy includes both packets and gaps between packets
  - Every container starts with a synchronization packet

Guaranteed streams to be aggregated

Input Ethernet Guaranteed (GST) streams:

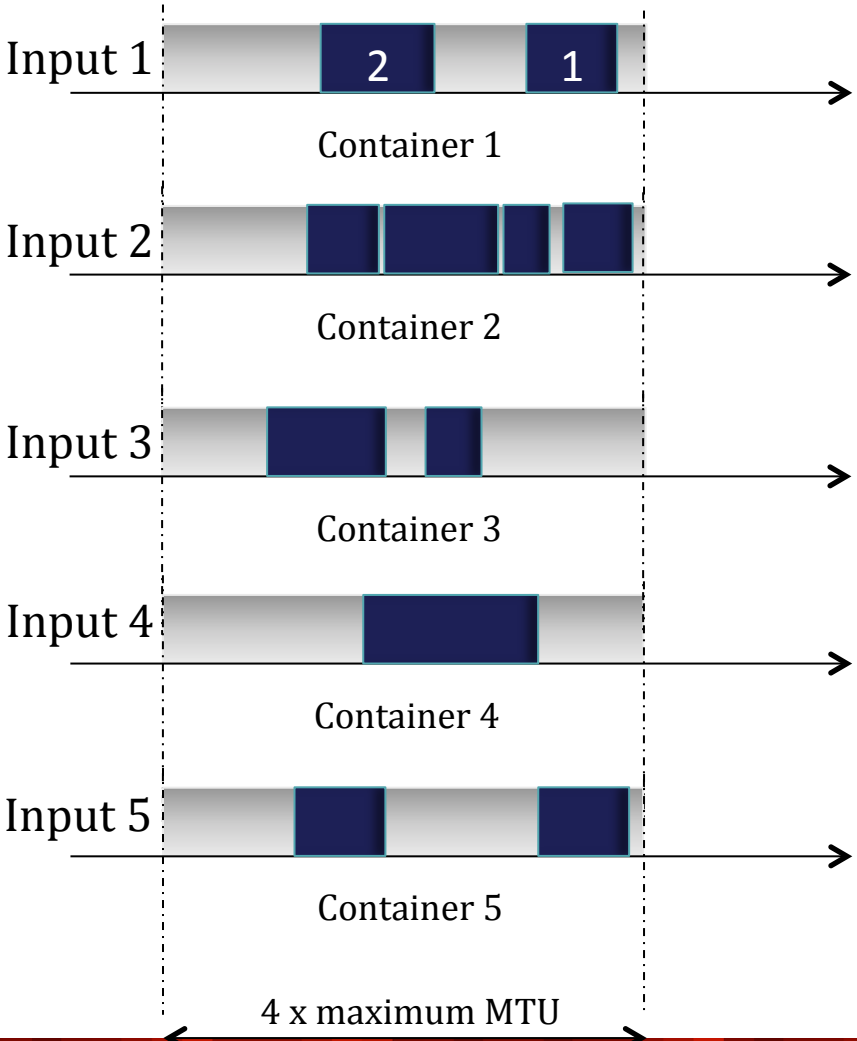

Fixed aggregation delay α :

Snapshot of all streams with packets and inter-packet lengths

Gaps preserved during aggregation

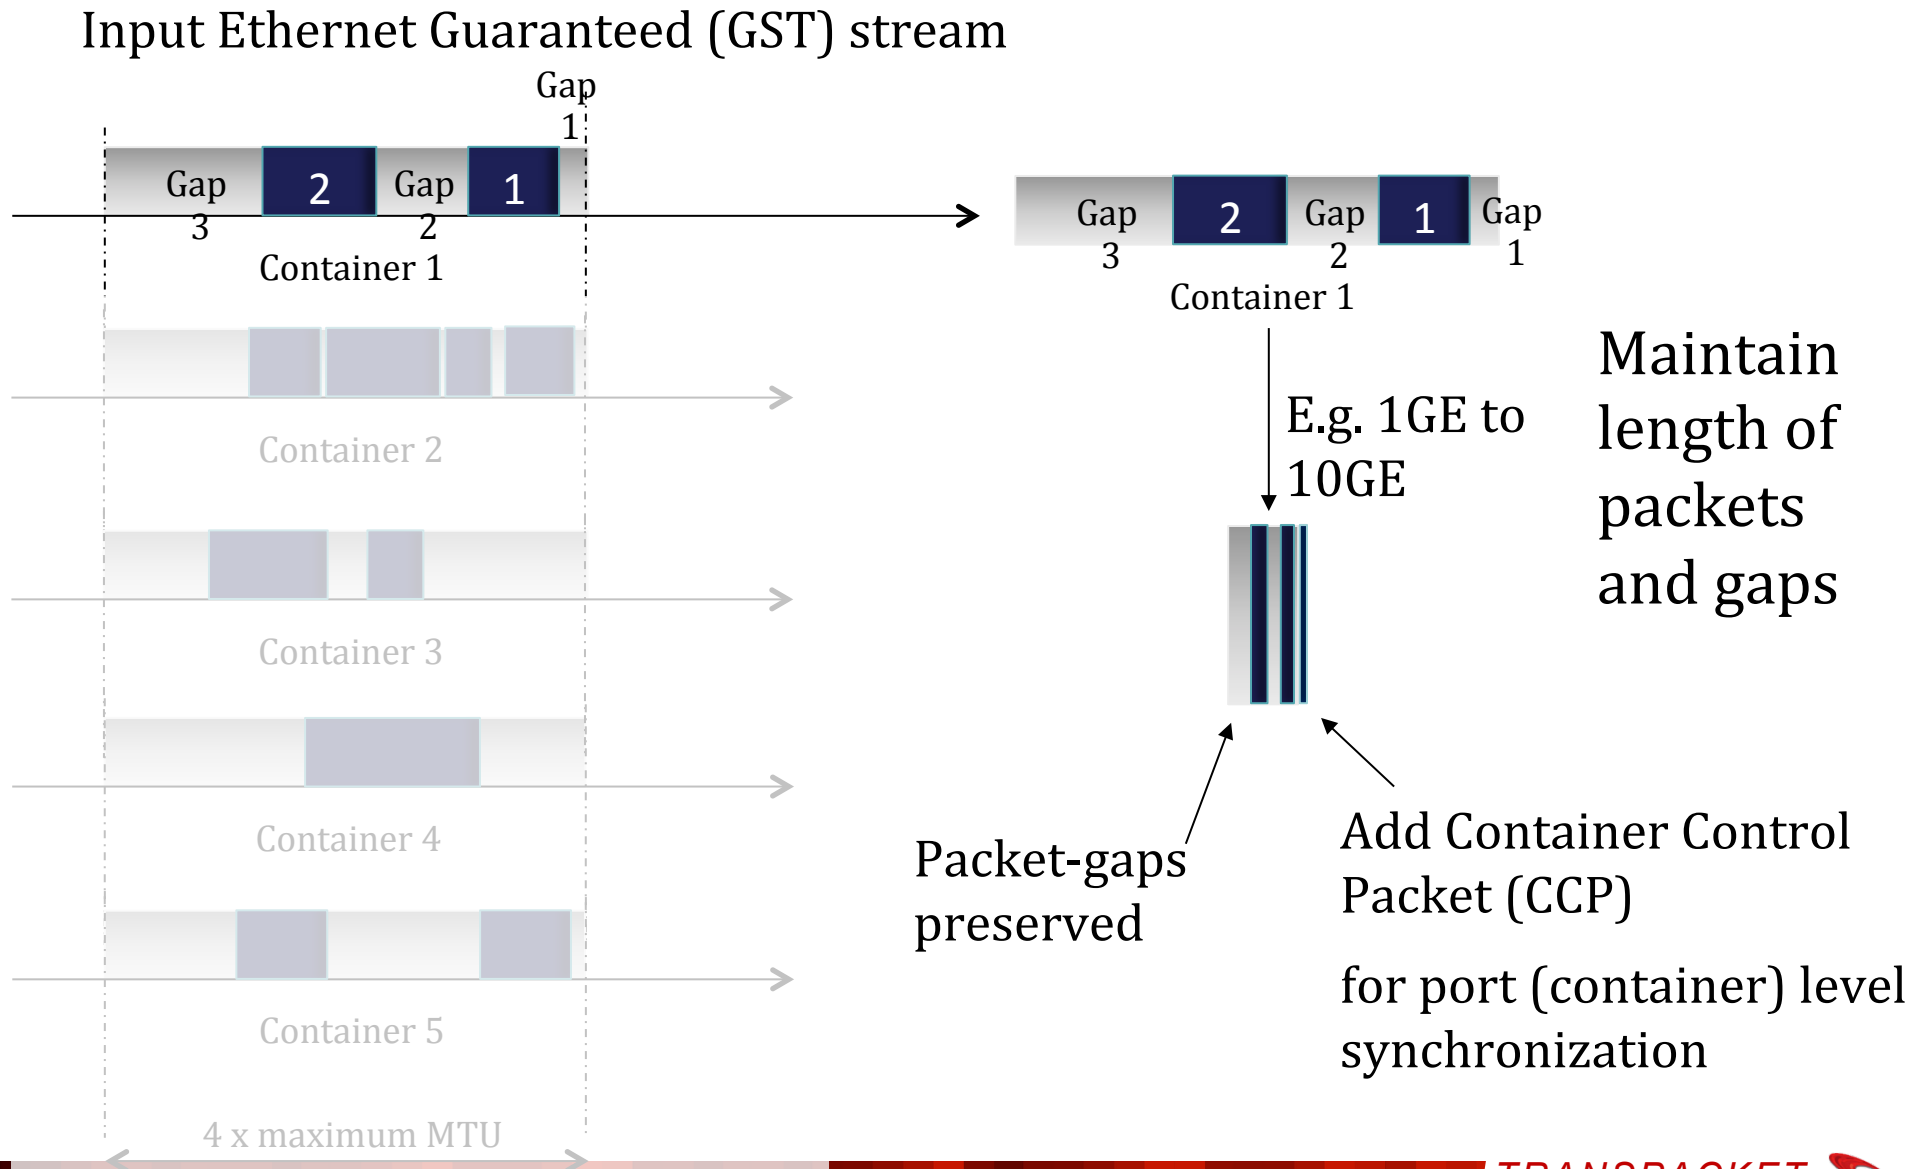

Aggregated guaranteed streams

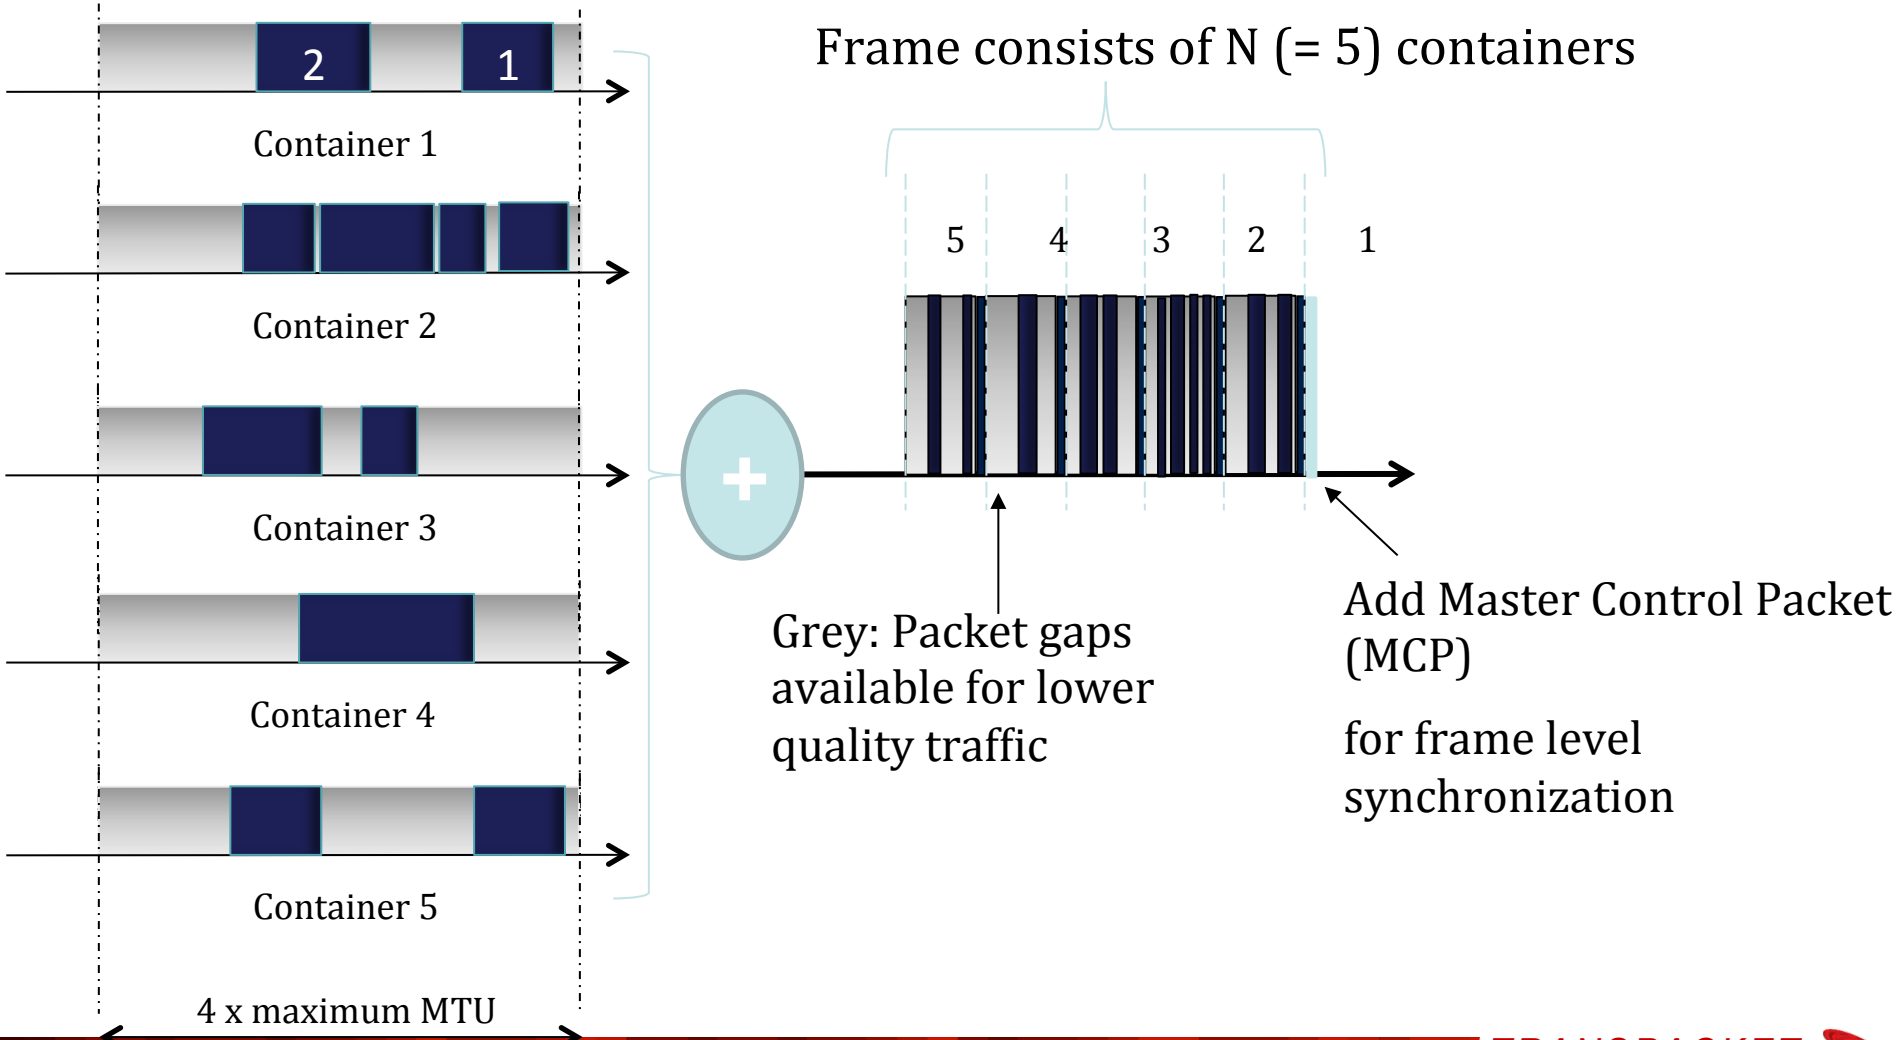

Adding packets into guaranteed streams using time-window approach

Aggregated GST streams

Aggregated GST streams with inserted
Lower quality statistically multiplexed (SM) streams

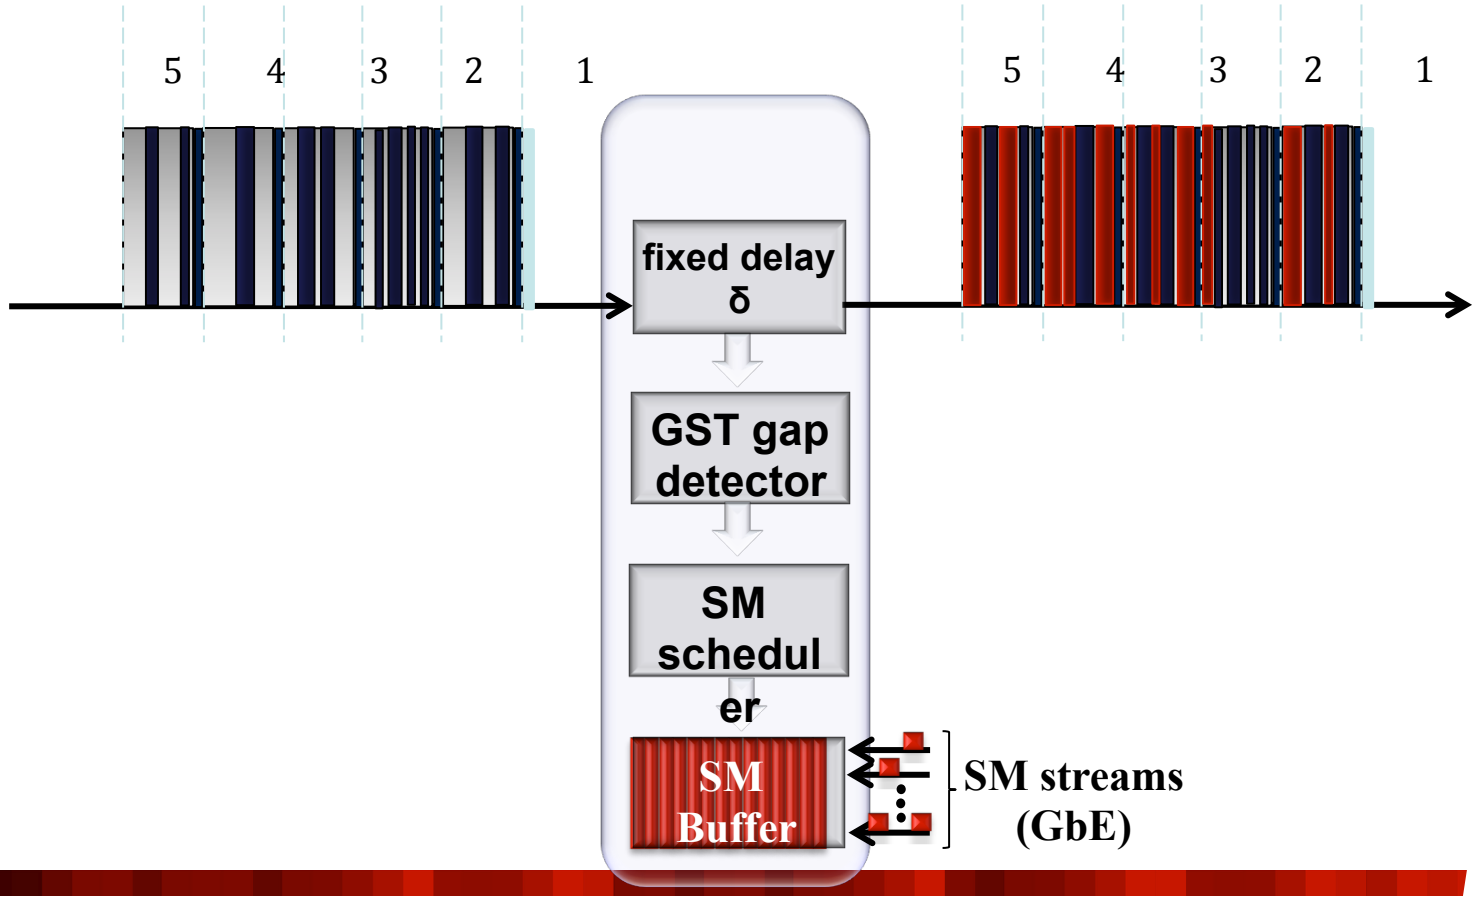

Performance of FUSION

Experimental results

Packet-jitter and latency sources

- Some packet jitter is unavoidable
  - Clock domain conversion, typ. < 10 ns at 10 Gb/s
  - Header processing (may be fixed)
  - Contention and buffering
    - Microseconds/milliseconds jitter
    - Jitter depends on traffic-load and scheduling mechanisms

10GE Guaranteed stream (GST)
2 X GE added using time-window (SM)

TransPacket H1 nodes

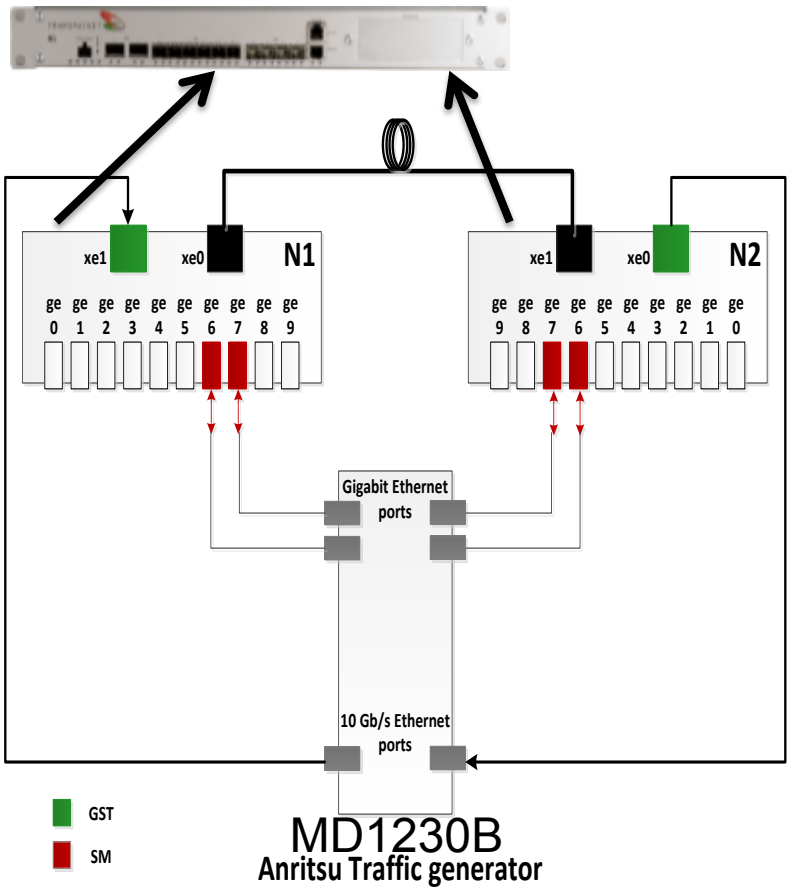

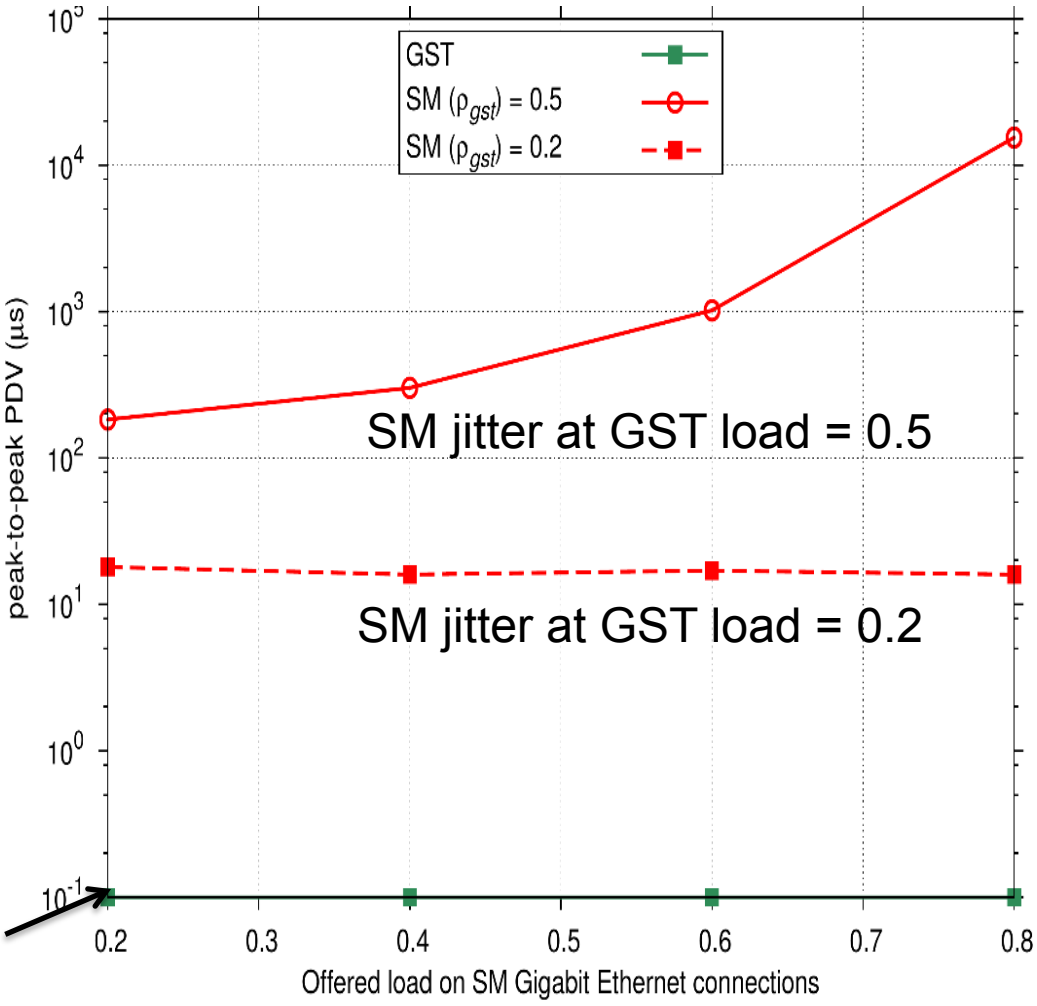

GST jitter independent of load

Packet-jitter GST 10GE stream for GST load = 0.5

- Y-axis: Frame count, X-axis: Latency intervals
- GST jitter < 50 ns and independent of SM load

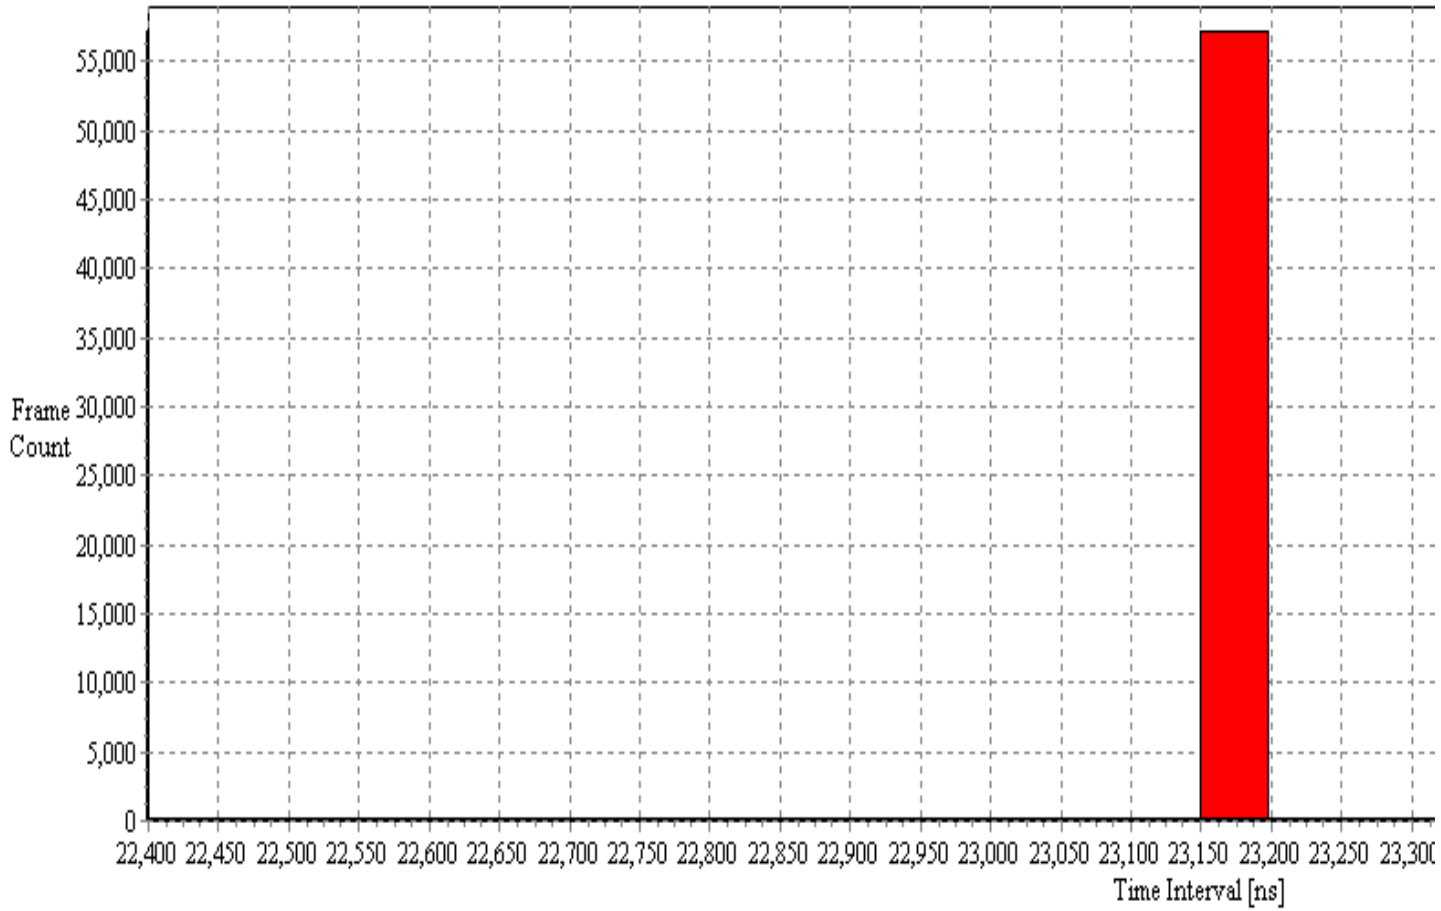

Packet-jitter for lower quality (SM) streams, GST load = 0.5

- Y-axis: Frame count, X-axis: Latency intervals
- SM jitter in the microsecond range

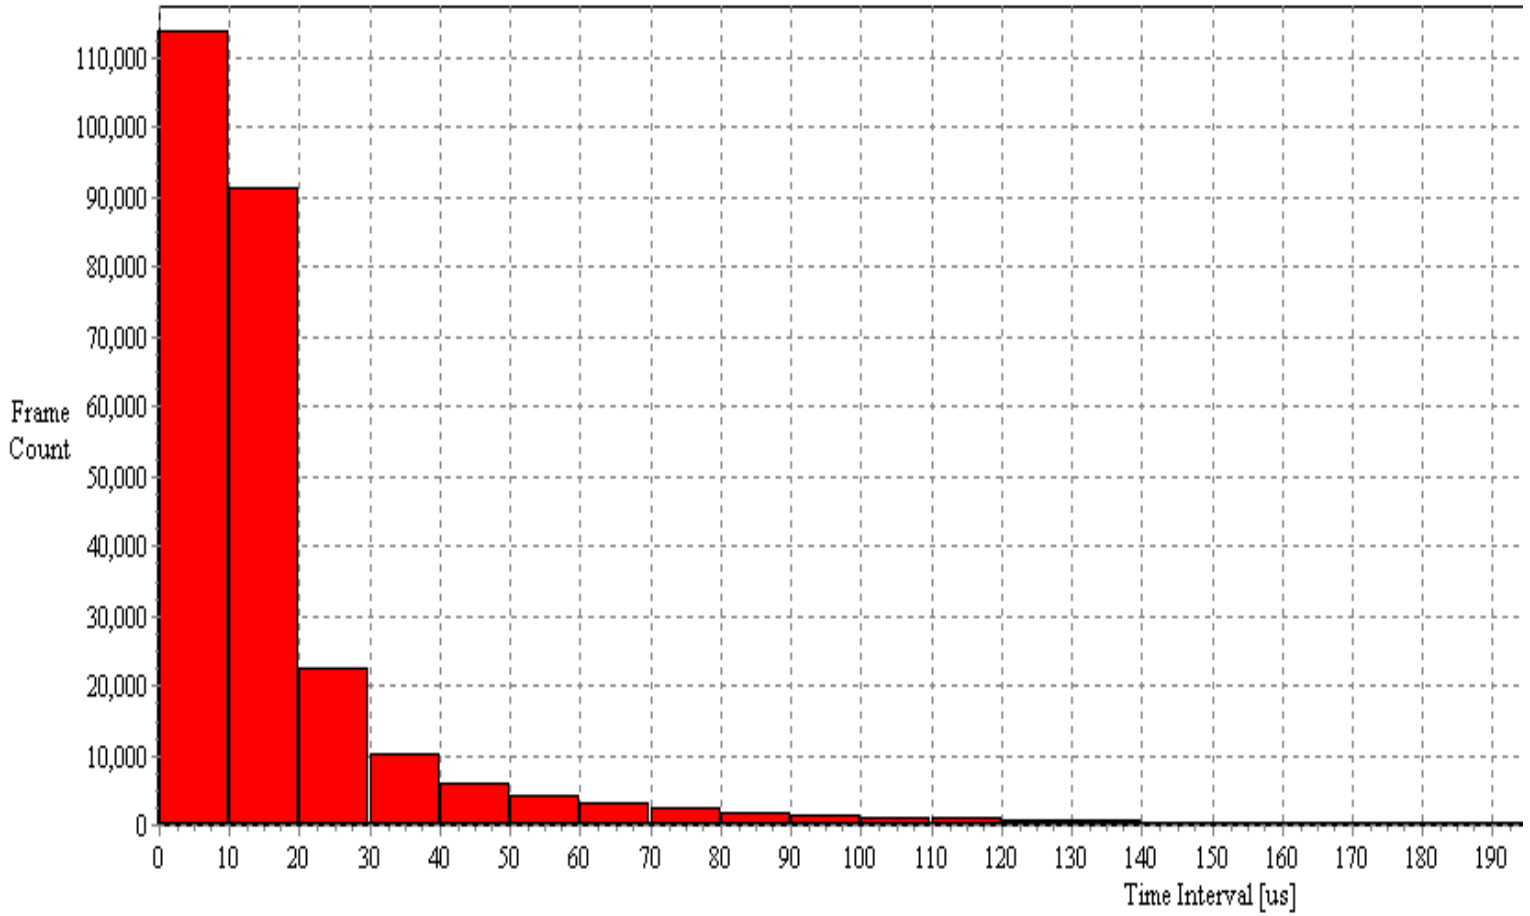

5 X Guaranteed streams (GST) + 2 X SM aggregation: P-Jitter for SM and GST

TransPacket H1 nodes

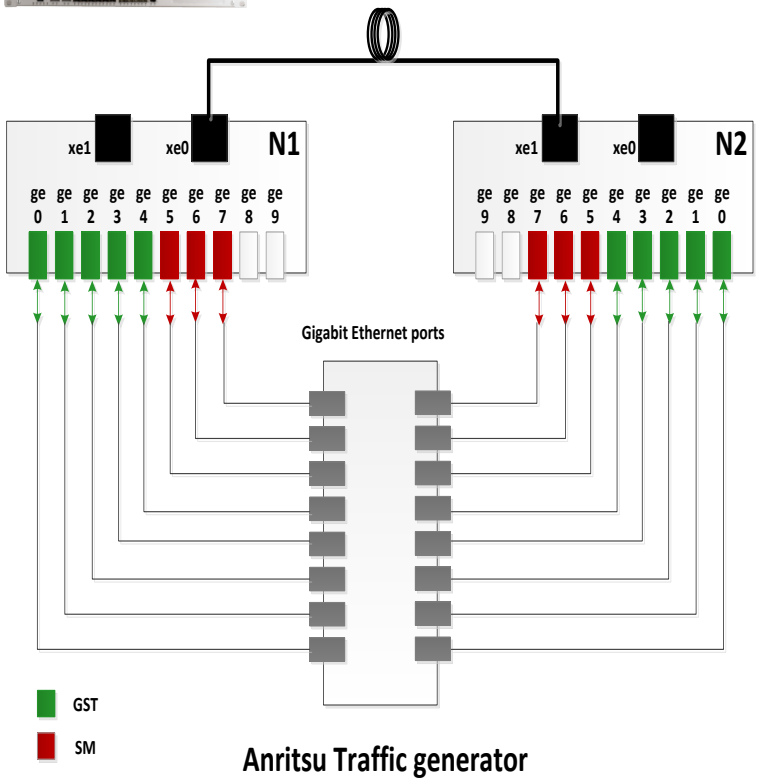

Anritsu Traffic generator

MD1230B

GST jitter < 250 ns

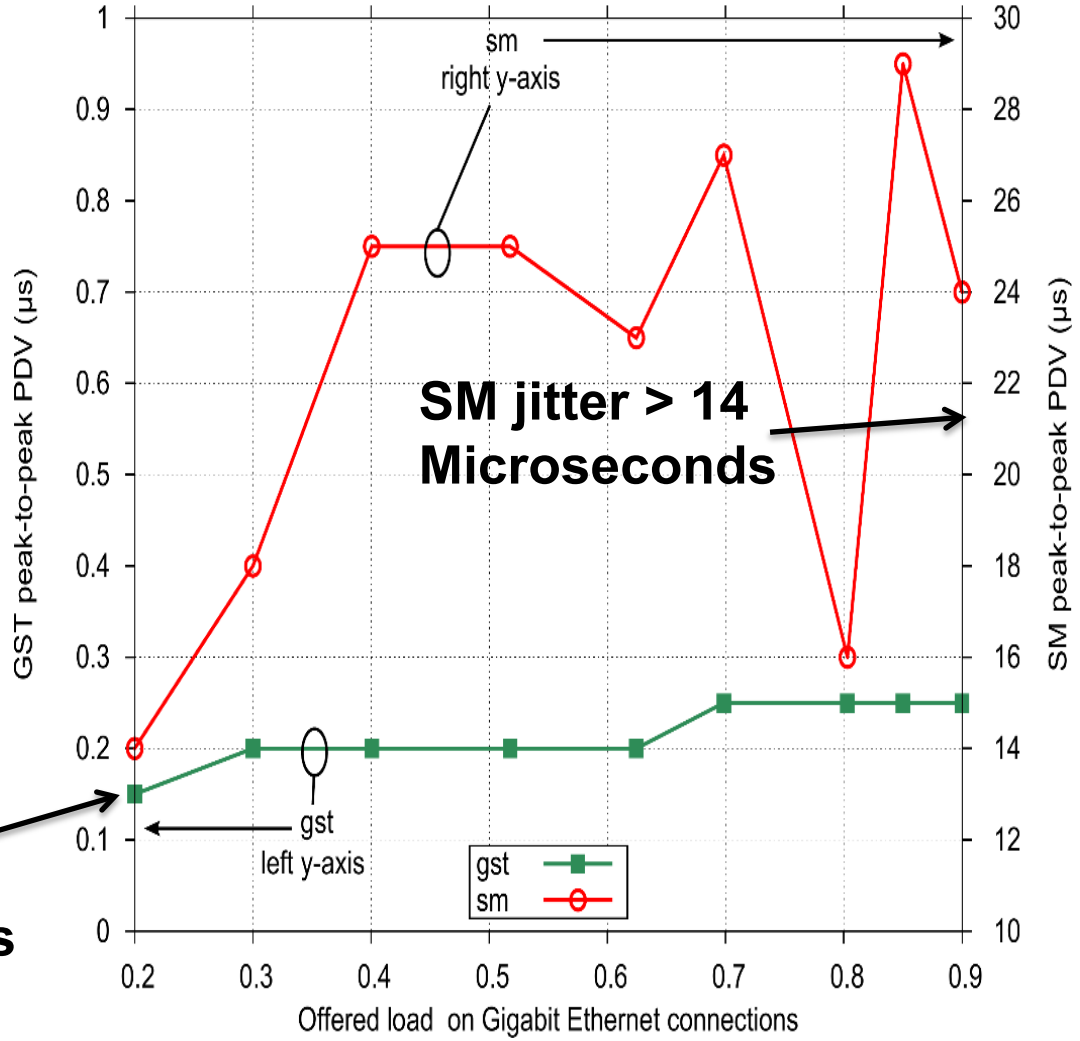

Summary and conclusion

- FUSION (IHON) Time-window minimizes packet-jitter
  - Virtual wavelength services
  - Mobile back-haul: IEEE1588 Sync accuracy depends on packet-jitter
  - Mobile front-haul (Radio over Ethernet, IEEE 1504.3), jitter < 50 ns achievable over several hops
- FUSION (IHON) guaranteed stream aggregation
  - Minimizes packet-jitter
  - Requires buffering at ingress (60 microsecond in experiment)

TRANSPACKET

FUSION NETWORKS