Frame Latency Issues in Multilane EPON

Glen Kramer, gkramer@broadcom.com

Introduction

NG-EPON

- This presentation explores the latency issues in a multilane channel-bonded PON where a single frame-based 100 Gb/s stream is carried over 4 wavelengths at 25 Gb/s.
- **Demultiplexor** splits a single serial 100 Gb/s ingress stream into 4 parallel 25 Gb/s egress streams
- **Multiplexor** combines 4 parallel 25 Gb/s ingress streams into a single 100 Gb/s egress stream.

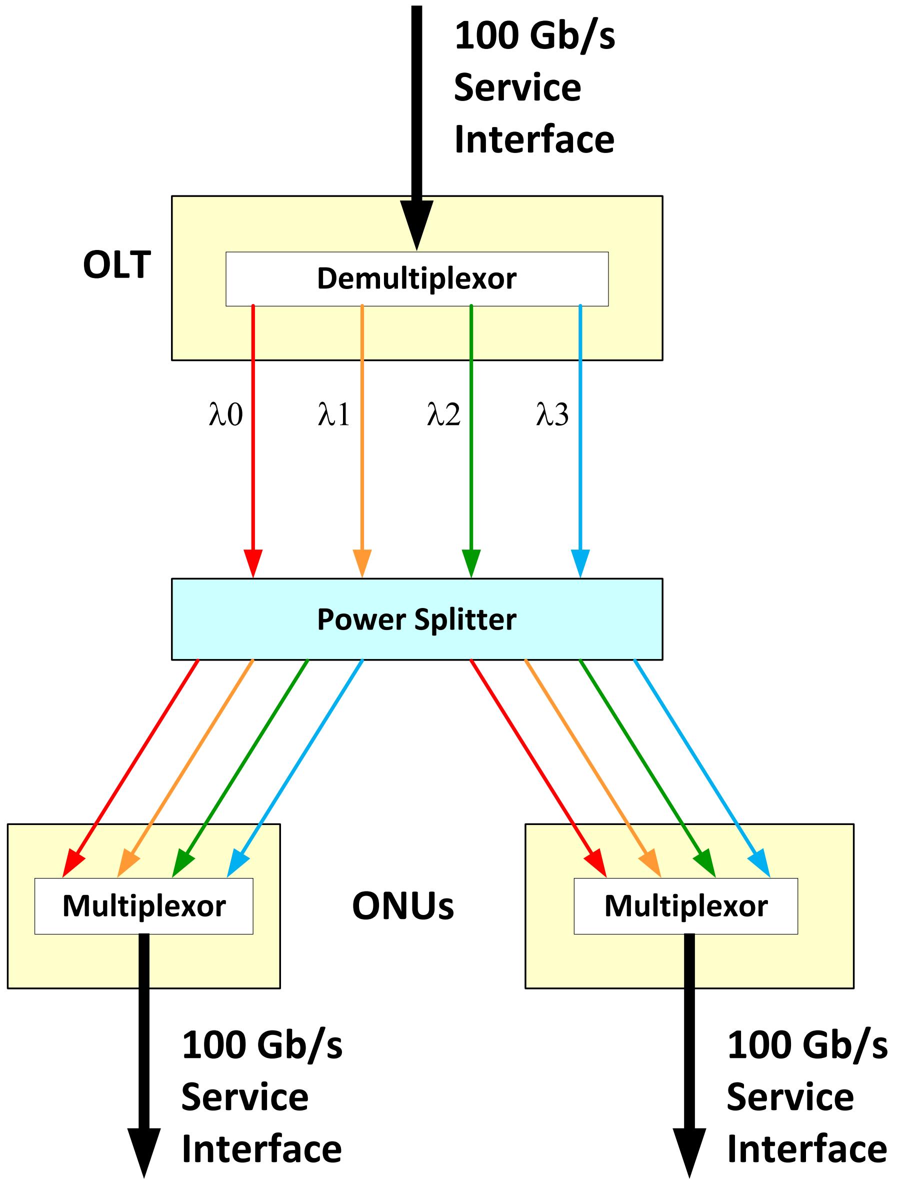

The Demultiplexor Model

- Frames arrive on a single MAC interface at 100 Gb/s
- Frames depart on 4 lanes at 25 Gb/s
- For each arriving frame, the Frame Distributor selects a lane with *earliest availability*
- When the first bit of a frame arrives, if a lane is available, transmission starts immediately (*cut-through* method)

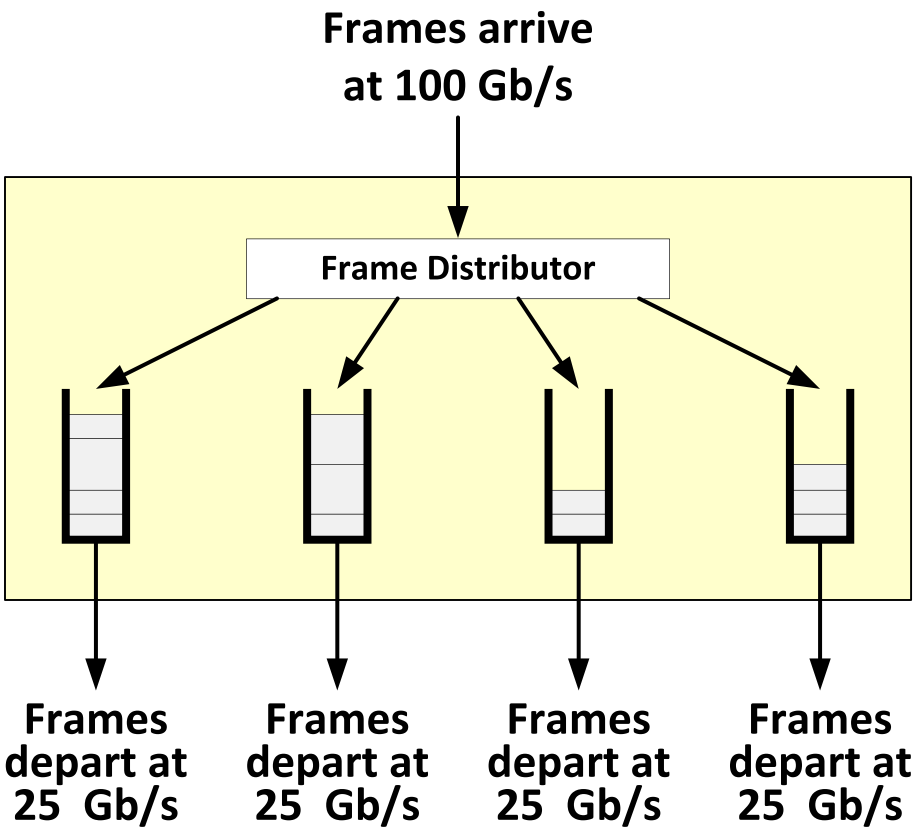

Demultiplexing Operation

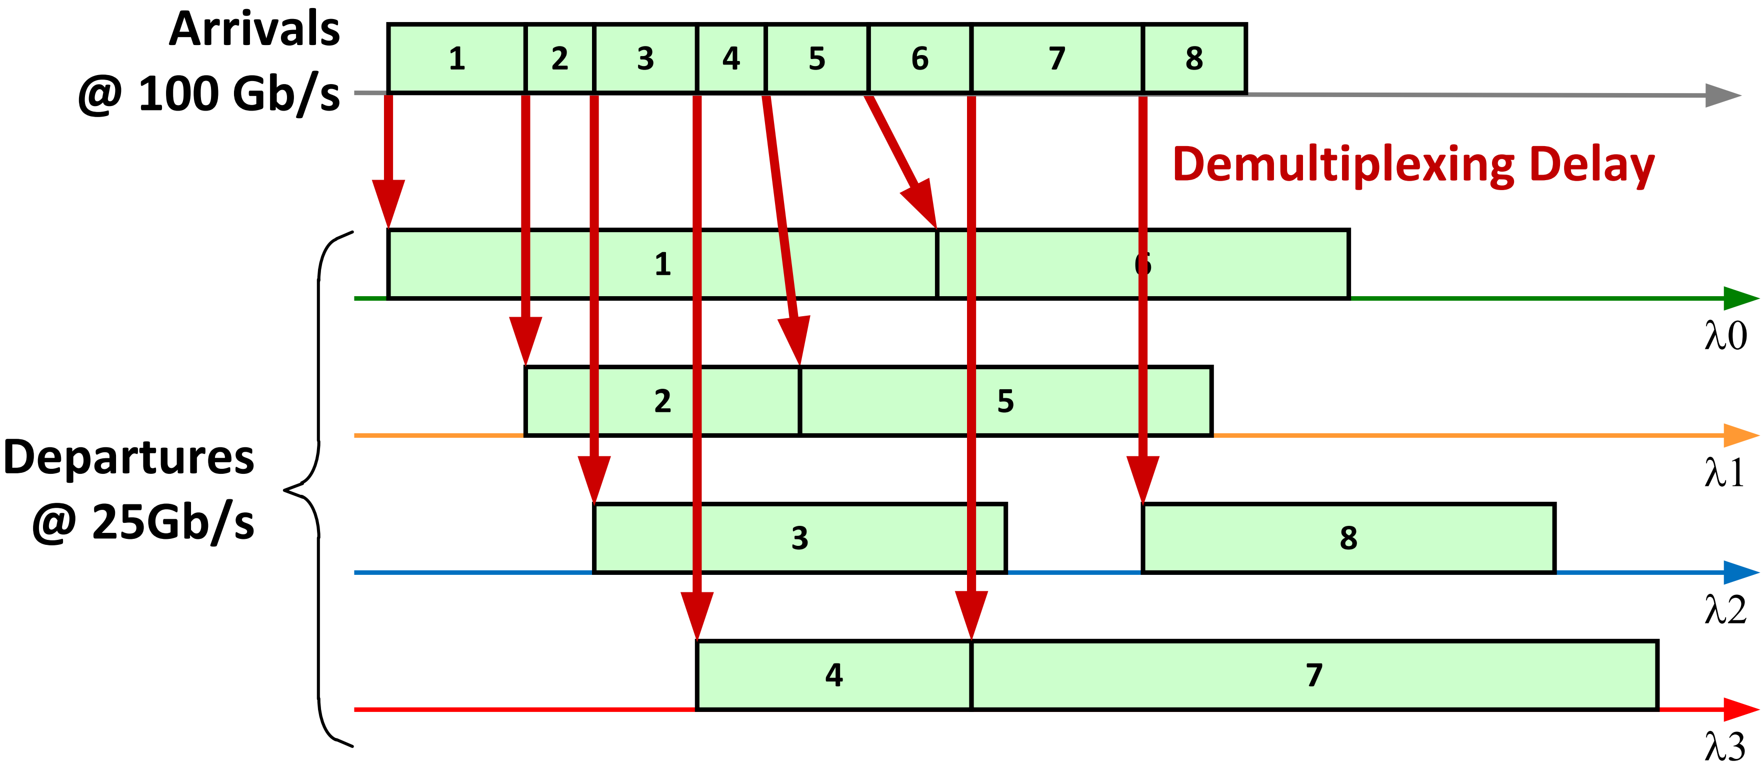

- Notice that frame 2 will finish the transmission before the frame 1, and frame 8 will finish its transmission before frame 7.
- Also notice that a gap is left in lane 2 because no frame is available at the ingress.

Smallest Delay = 0

- If a REPORT MPDPDU arrives, when a lane is available, it will be sent with zero delay

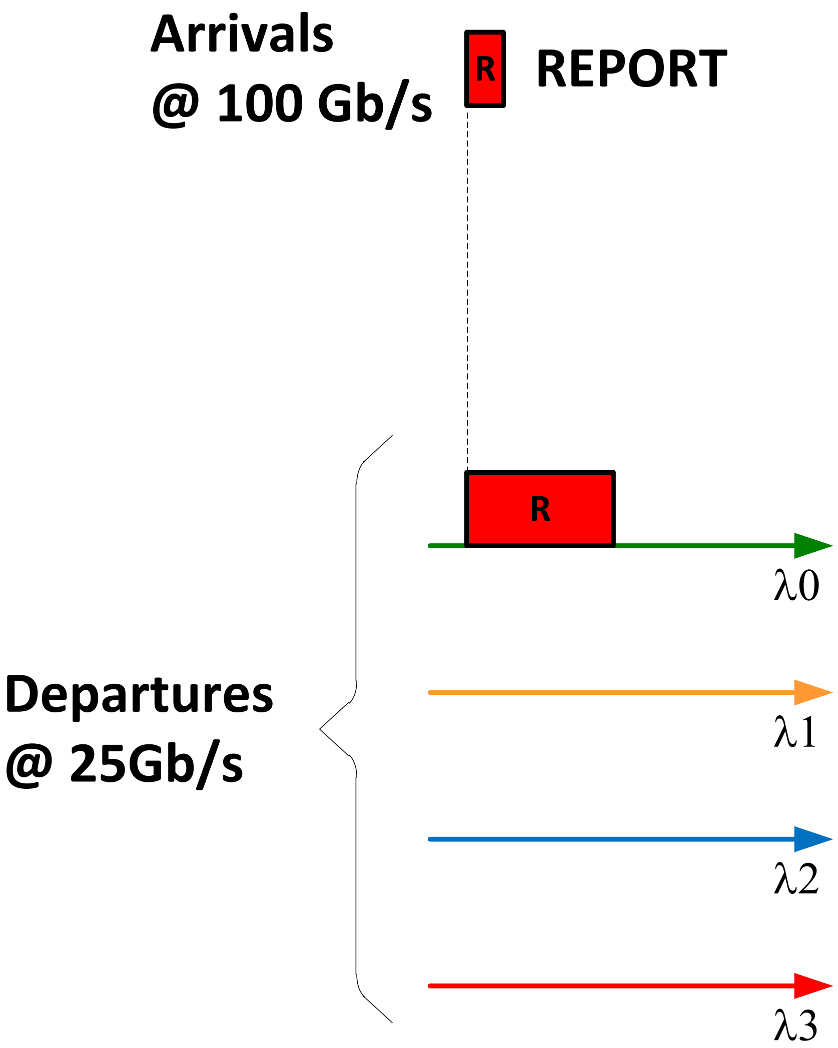

Unlucky case (illustration is on the next slide)

- Now consider a case where all lanes are also initially available, but REPORT follows 4 data frames of sizes X1, X2, X3, and X4, such that

  $X2 = \frac{3}{4} X1$

  $X3 = \frac{3}{4} X2$

  $X4 = \frac{3}{4} X3$

- In this situation, all data frames will start transmissions immediately upon arrivals of their first bits, and all will finish at the same time.
- The REPORT will arrive at time $X1+X2+X3+X4 \approx 2.7 \times X1$, but will find no lanes available until the time $4 \times X1$. So, the REPORT will experience a delay equal to 1.3 times the transmission time of frame X1 @100 Gb/s.
  - If X1 is a 2000-byte frame, REPORT delay will be 208 ns.
  - If X1 is a 9000-byte frame, REPORT delay will be 936 ns.
- According to simulations, this is not the worst case scenario

Unlucky case

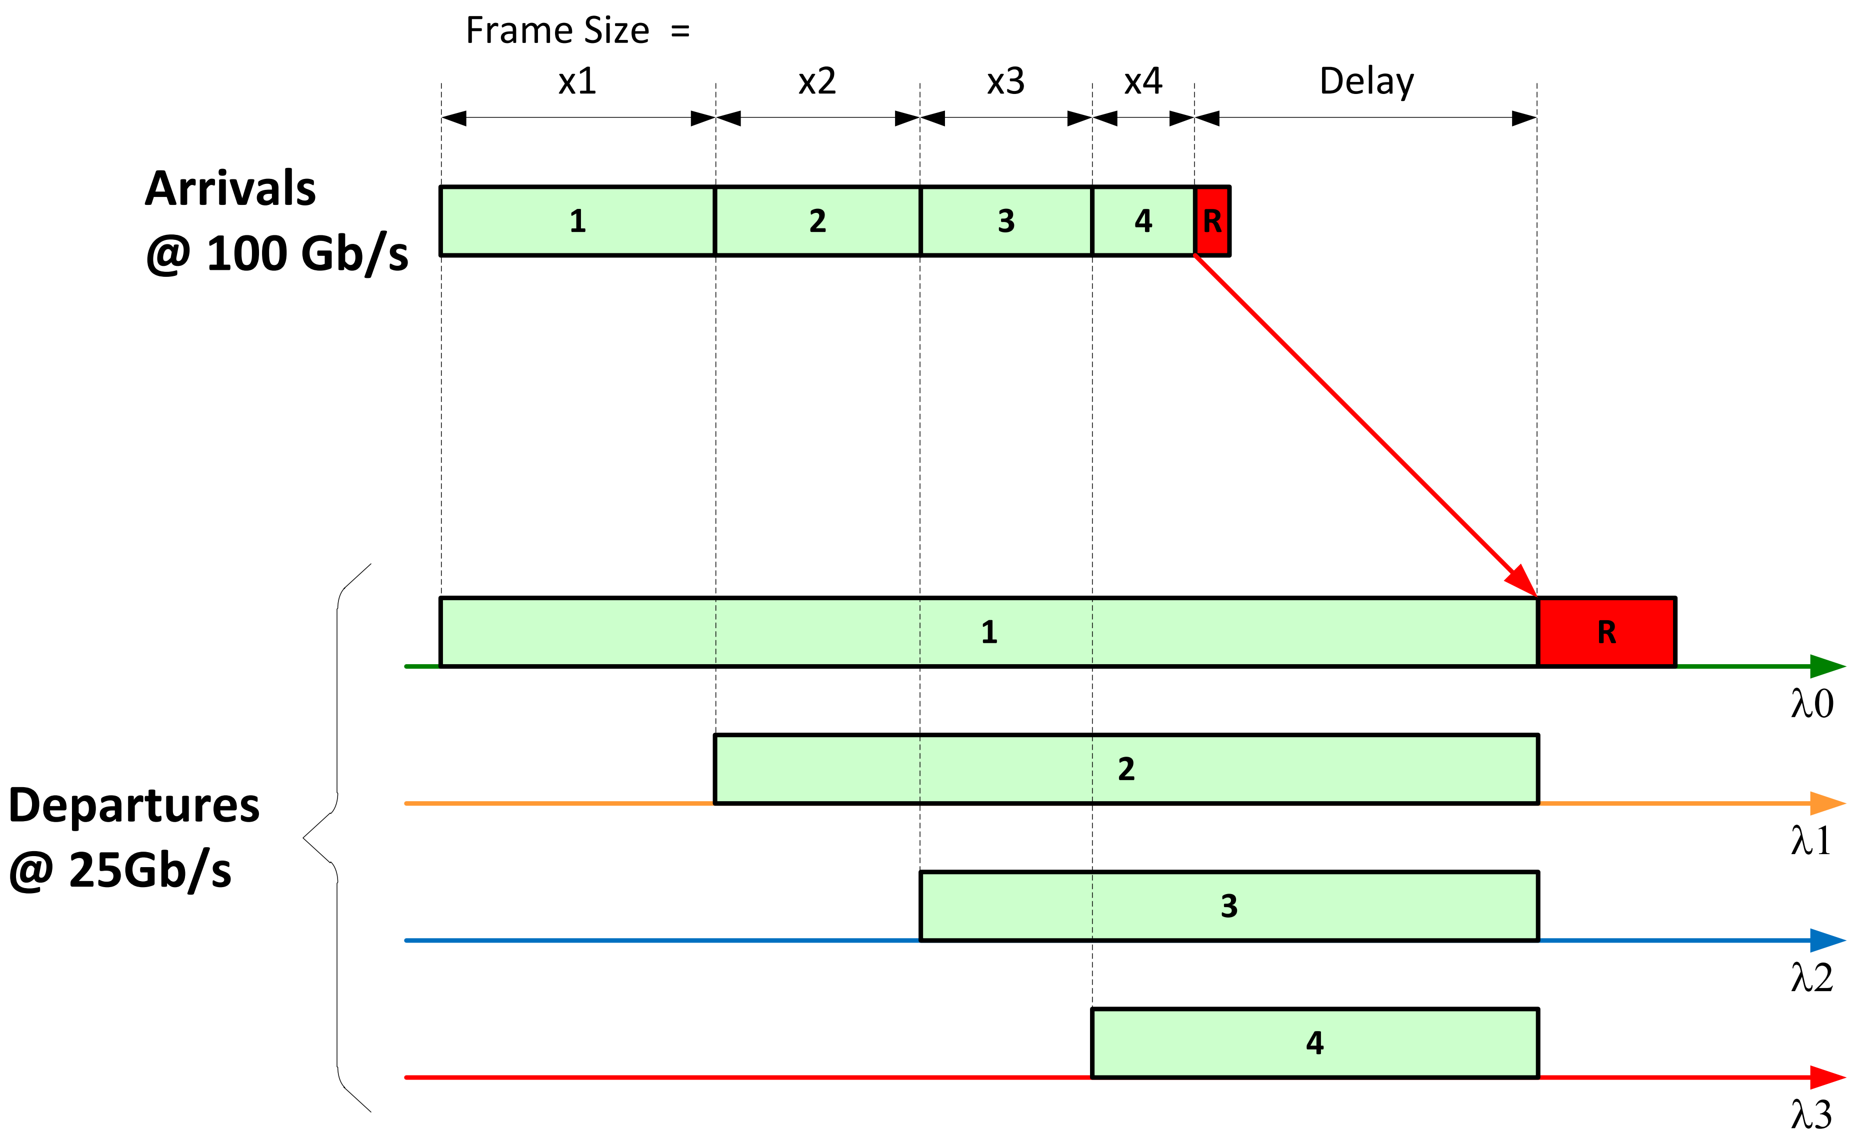

The Multiplexor Model

- Frames arrive on 4 lanes at 25 Gb/s
- Frames depart on a single MAC interface at 100 Gb/s
- An entire frame must be received at 25 Gb/s before it can be transmitted at 100 Gb/s (*store-and-forward* method)
- Frames may arrive on 4 lanes out of order. The Frame Combiner (Serializer) must restore the original frame order.

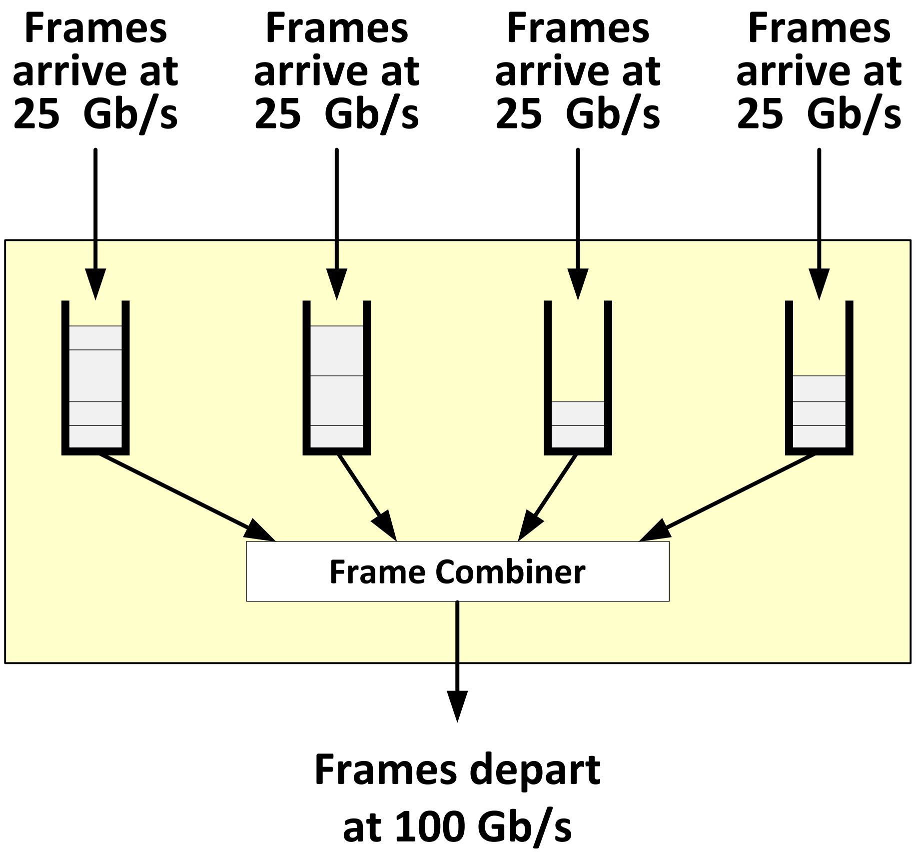

Multiplexing Operation

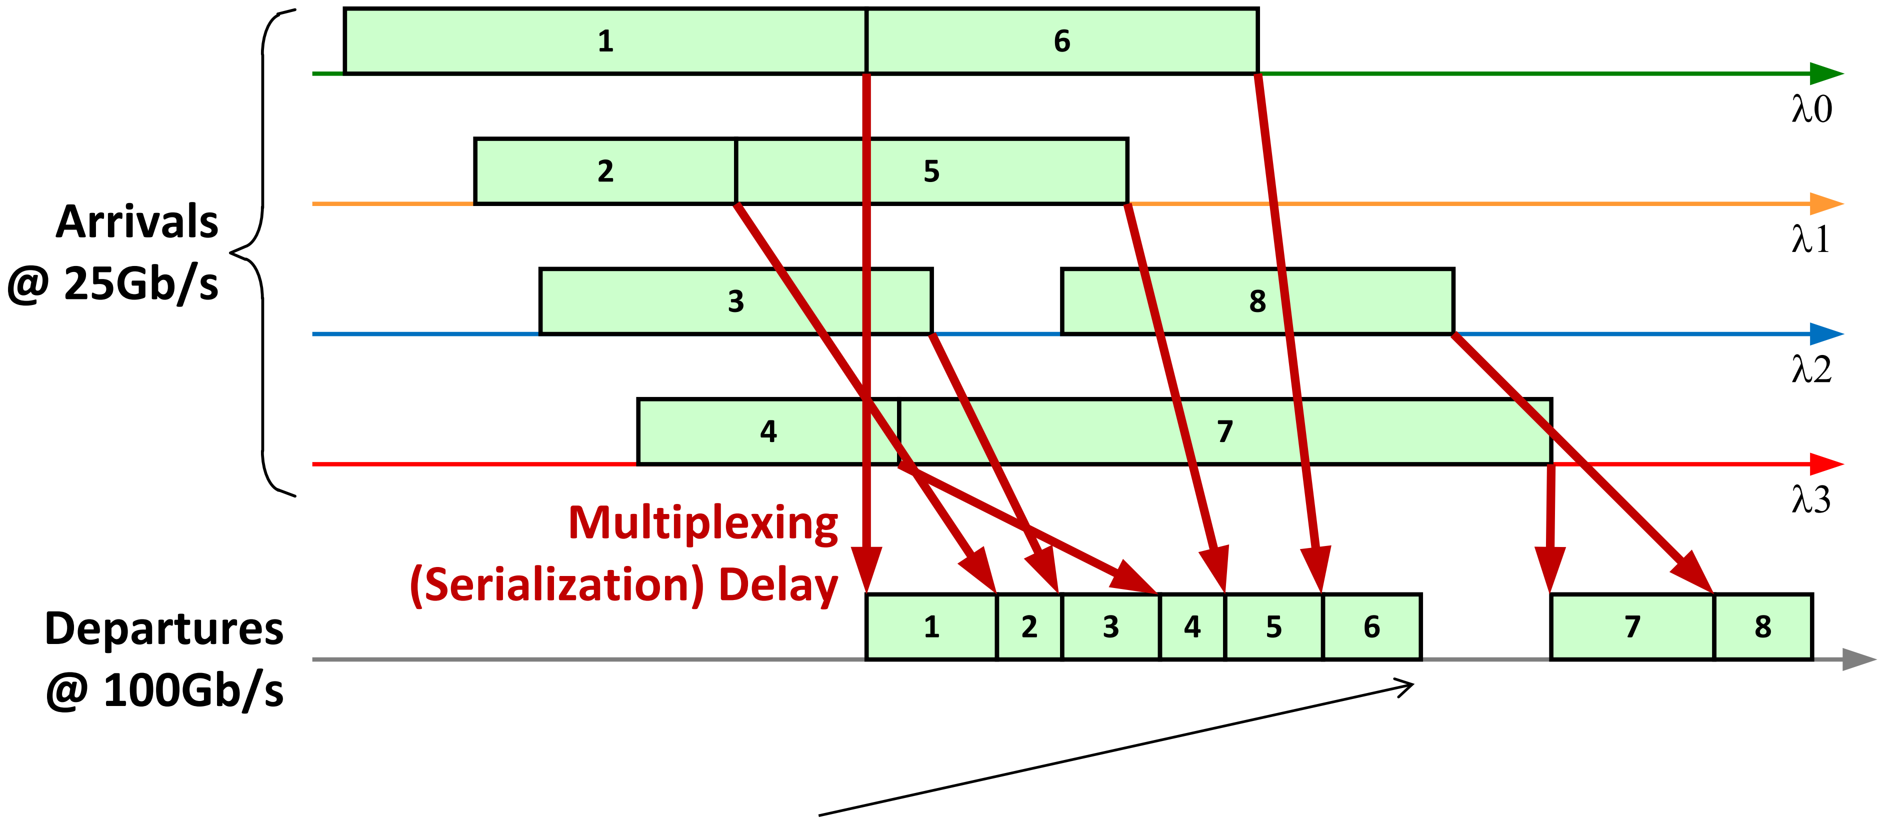

- Note the presence of gaps in the egress stream due to frame unavailability. No such gaps were present in the 100 Gb/s ingress stream in the Demultiplexor.
- Gaps in the egress stream mean that some queue somewhere in a galaxy far far away gets longer and longer.

End-to-End Delay

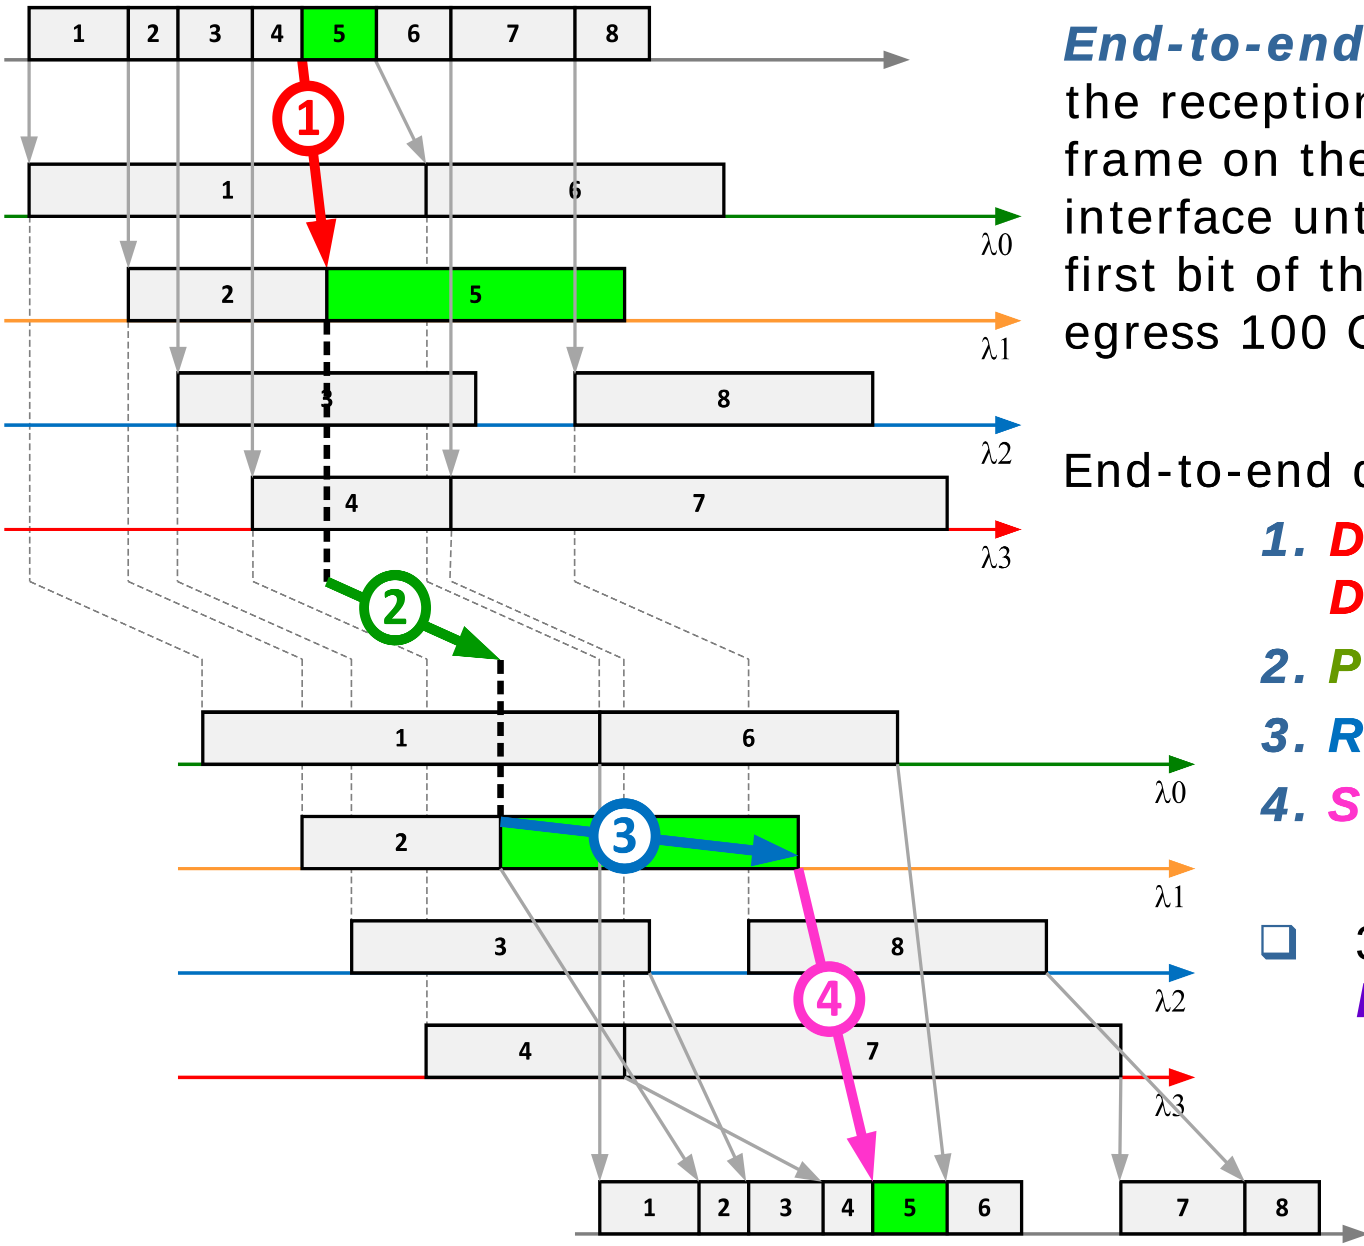

End-to-end delay is time from the reception of first bit of a frame on the ingress 100 Gb/s interface until transmission of the first bit of this frame on the egress 100 Gb/s interface.

End-to-end delay consists of

1. **Demultiplexing Delay**
2. **Propagation Delay**
3. **Reception Delay**
4. **Serialization Delay**

- 3 and 4 together is ***Multiplexing Delay***

Analysis

- Model is built in Excel
- 10K frames
- Frame sizes: random uniform in [64…2000] (bytes)
- Frame overhead: 20 bytes (12 IPG + 8 preamble)
- Propagation delay is assumed to be 0 ns.

- Excel model is available at …

Reception & Serialization Delays

- Reception and Serialization delays are inversely correlated
  - The more time it takes to receive a frame, the less time the frame will need to wait to get on the serial egress channel.
- Thus, we can expect that the multiplexing delay will have smaller variability compared to either reception delay or serialization delay

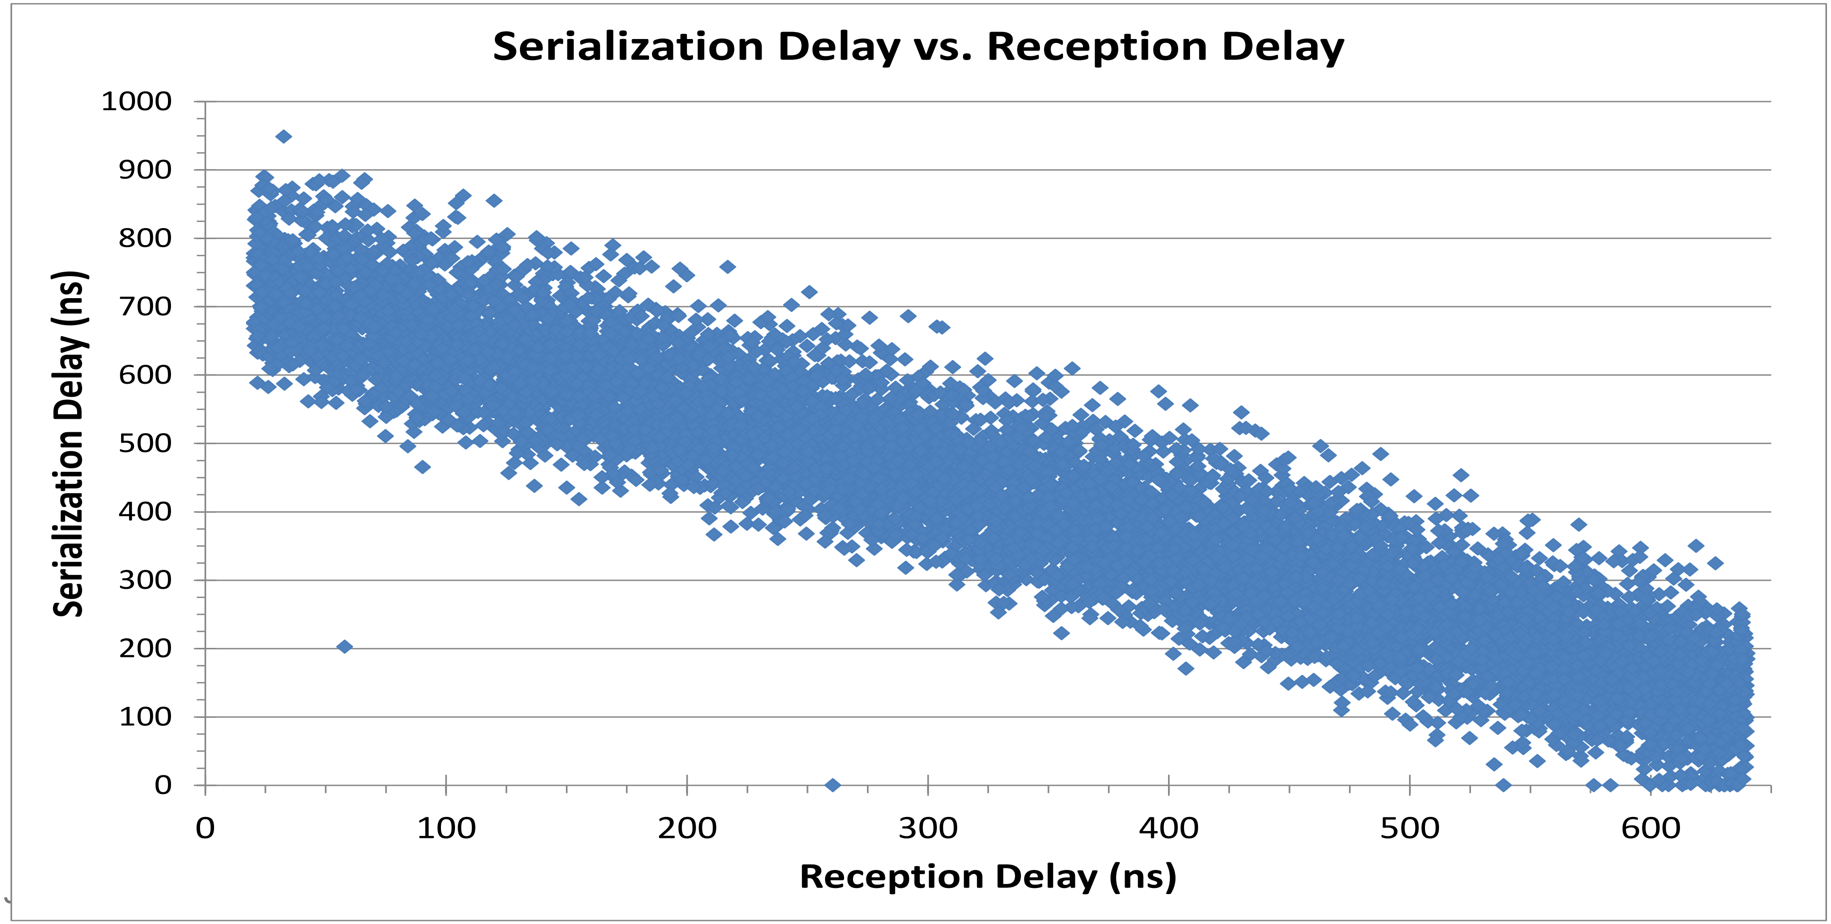

Multiplexing Delay

NG-EPON

Multiplexing Delay = Reception Delay + Serialization Delay

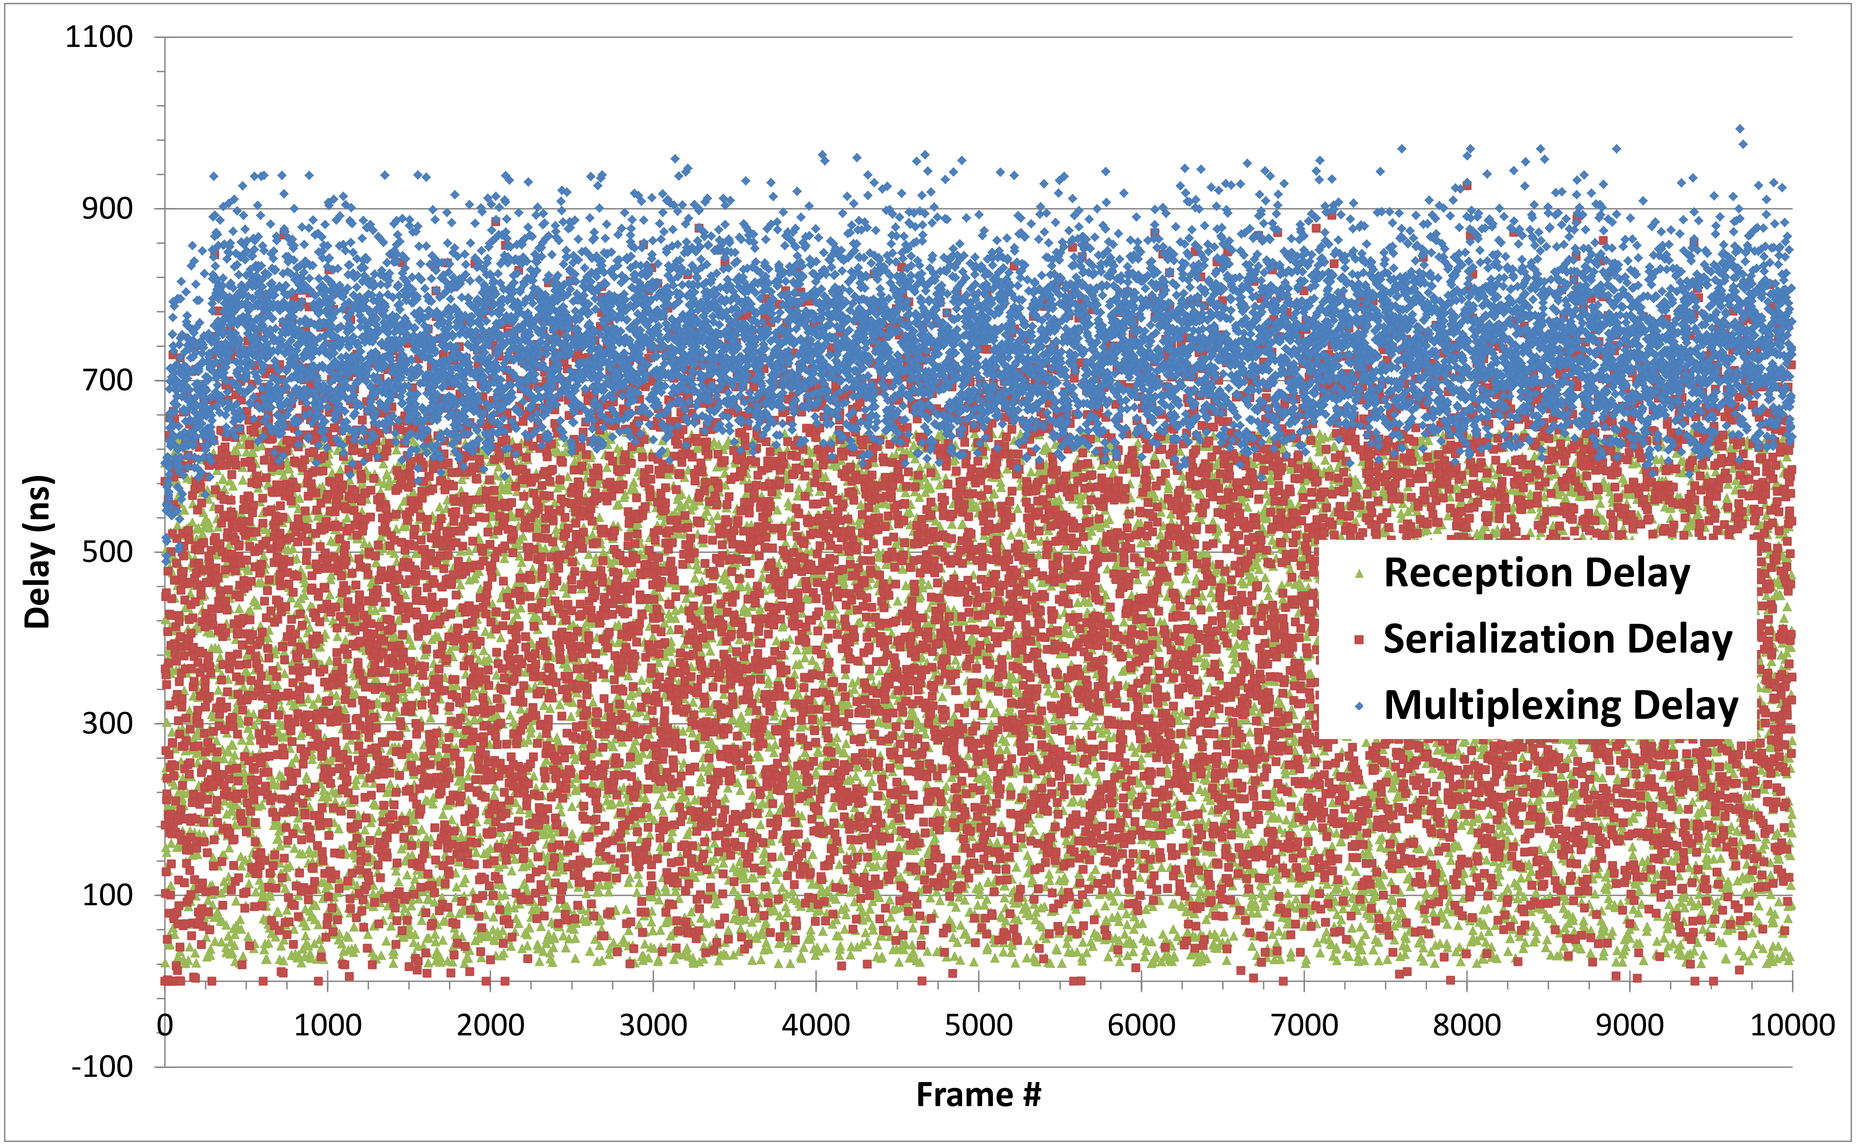

Demultiplexing & Multiplexing Delays

- Demultiplexing and Multiplexing delays are inversely correlated
  - The more time a frame waits in the OLT, the less time it will wait in the ONU, and vice versa.
- We can expect that the End-to-End Delay will have smaller variability than either the demultiplexing delay or multiplexing delay

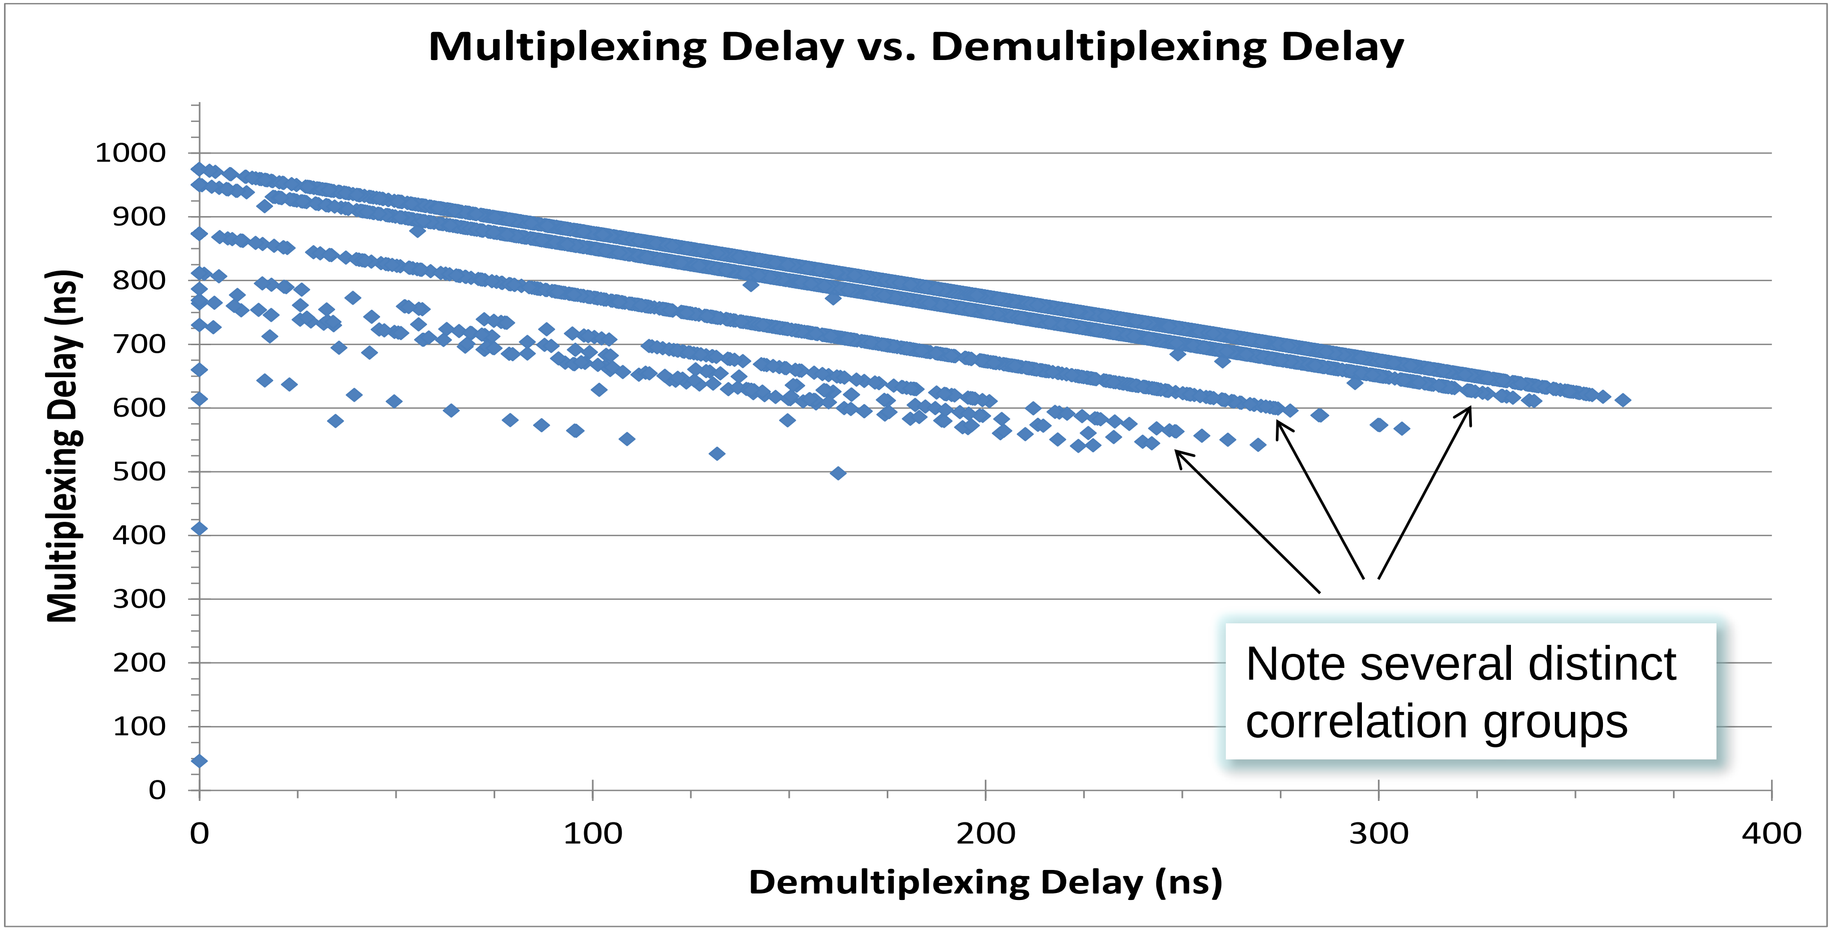

End-to-End Delay Analysis

- End-to-End Delay has low variability (constant within a correlation group)

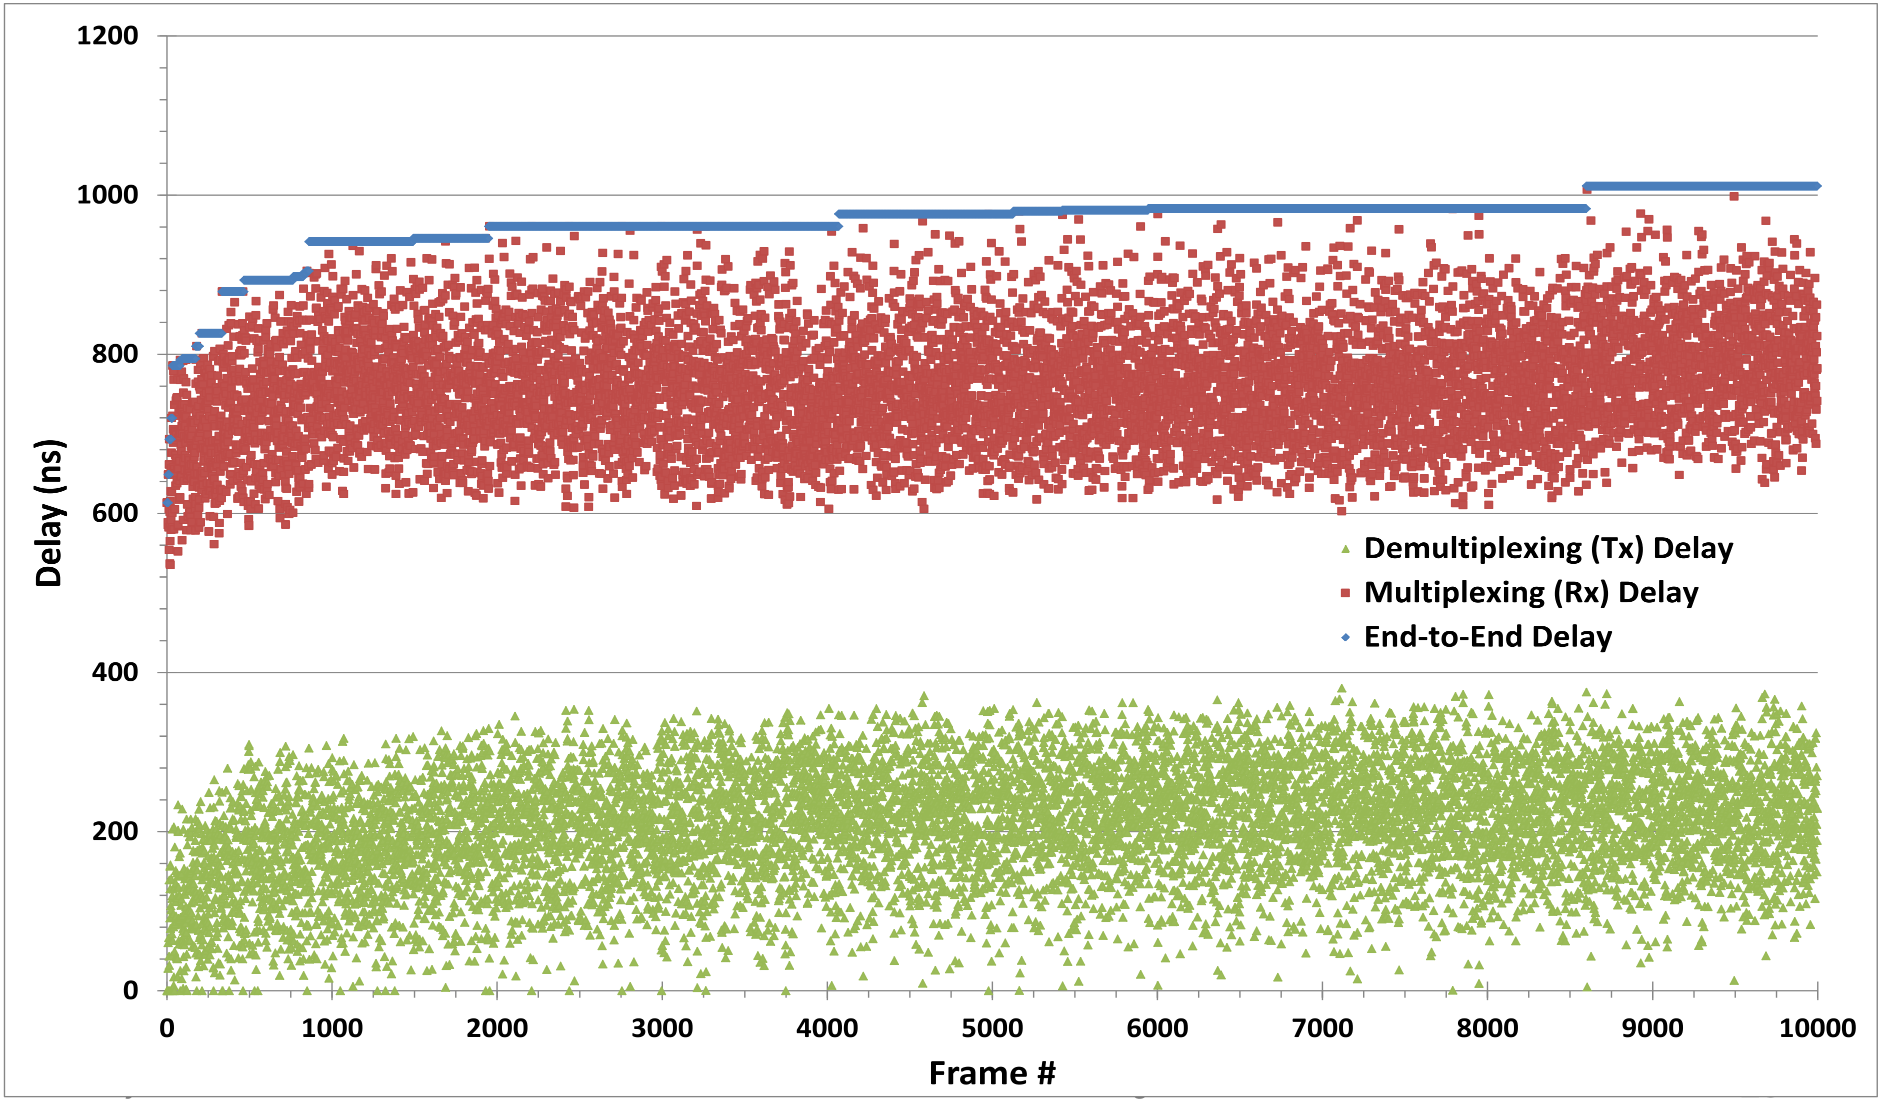

End-to-End Delay Jumps

- Every time the egress 100 Gb/s stream has a transmission gap, the end-to-end delay jumps to a new high

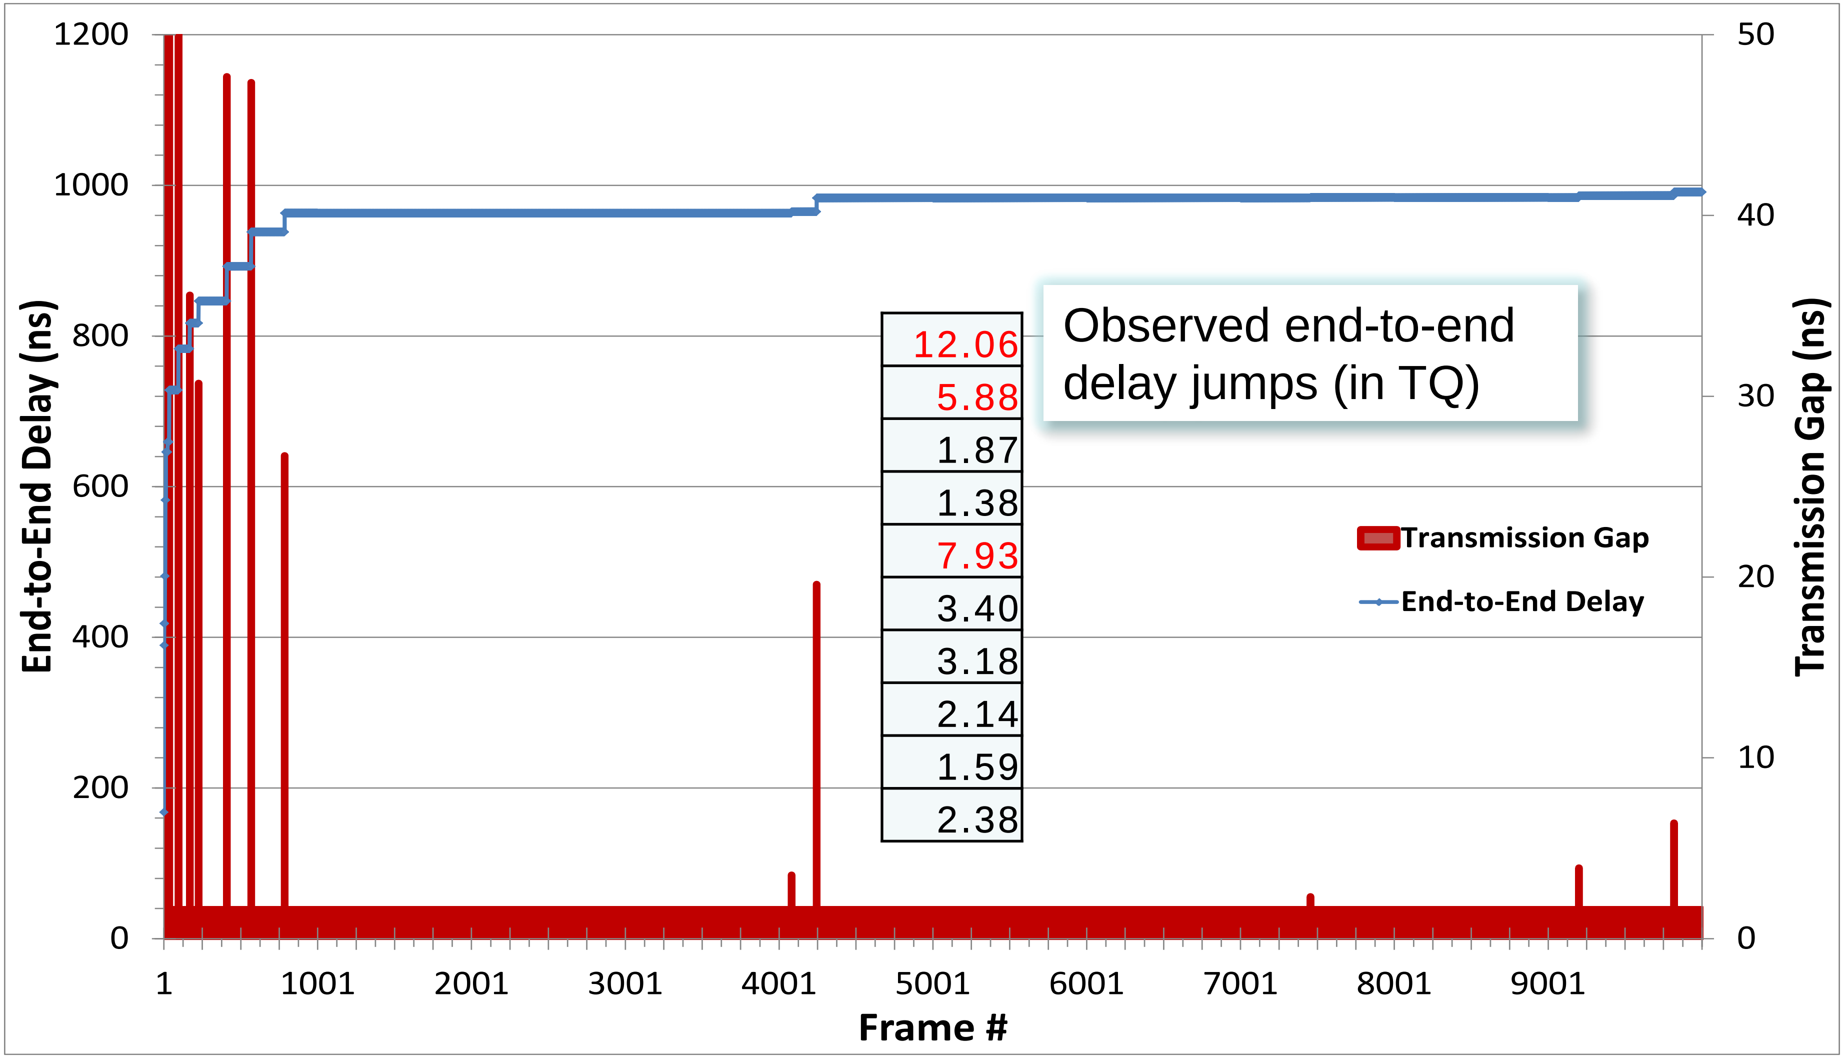

Key Observations

NG-EPON

- The End-to-End delay variability is much lower than the delay variability at either end (in the OLT or ONU) alone.
- End-to-End delay depends on maximum frame size
  - 2000-byte max frame: ~ 1μs
  - 9000-byte max frame: ~ 4.5μs
- This delay is acceptable to some user traffic.
  - Delay variability through EPON is still dominated by queuing delay while a frame waits for upstream timeslot (milliseconds).
  - Impact on 802.3AS or 1588?
- This delay variability breaks MPCP
  - Some jumps in End-to-End Delay variability will cause ONU deregistration due to timestamp drift
  - The Demultiplexing Delay variability will lead to upstream burst collisions
  - MPCP issues are explained in presentation “*Options for placing the channel bonding sublayer*”

Thank You