ACM MHV 2023

RMTS: A Real-time Media Transport Stack Based on Commercial Off-the-shelf Hardware

Ping Yu, Yi Wang, Ming Li, Jianxin Du, and Raul Diaz

Presented by: Raul Diaz, Sr. Principal Engineer
Network and Edge Group, Office of the CTO, Software Video Architecture

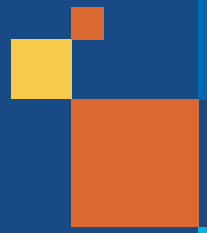

intel®

Media production transformation

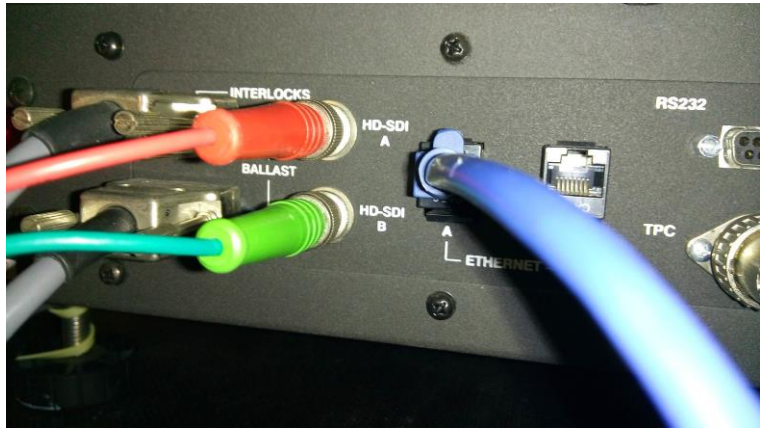

SDI

- 1989: SMPTE 259, 292, …
- Synchronous network (SDI)
- Standards tied to bit fixed bit rates
- Difficult to scale
- Expensive

FROM SDI networks and custom media processing equipment

TO IP-connected general purpose computer servers

ENABLES Virtualized media production

IP

- 2017: SMPTE ST 2022, ST 2110, …
- Asynchronous network (Ethernet, IP)
- Any bit rate, any media protocol
- Easy to scale
- Inexpensive

SMPTE ST 2110

- "The SMPTE ST 2110 standards suite specifies the carriage, synchronization, and description of separate elementary essence streams over IP…"
  - https://www.smpte.org/smpte-st-2110-faq
- RMTS is compliant with the full ST 2110 standard
- Paper focus: ST 2110-21 – most challenging part of ST 2110 for real-time video packet transport

Standard	Description
ST 2110-10	System Timing and Definitions
ST 2110-20	Uncompressed Active Video
ST 2110-21	**Packet pacing and Delivery Timing for Video**
ST 2110-22	Compressed Video
ST 2110-30	PCM Digital Audio
ST 2110-31	AES3 Transport Digital Audio
ST 2110-40	Ancillary Data
ST 2022-7	Seamless Protection Switching of SMPTE ST 2022 IP Datagrams
RFC 826	Address Resolution Protocol (ARP)
RFC 3376, RFC 2236	Internet Group Management Protocol (IGMP)
IEEE 1588	Precision Time Protocol, Time stamping (synchronization)

SMPTE ST 2110-21

Traffic Shaping and Delivery Timing for Video

- TR_{OFFSET}
  - Time at which the first packet of video data in the frame must be transmitted
  - Example: HD 1920x1080, T_{OFFSET} = 666µs
- T_{RS}
  - Time interval between video packets
  - Example: HD 1920x1080, 60FPS, 10bit color, T_{RS} = 3.57µs (For comparison UHD 3840x2160 T_{RS} = 579ns!)

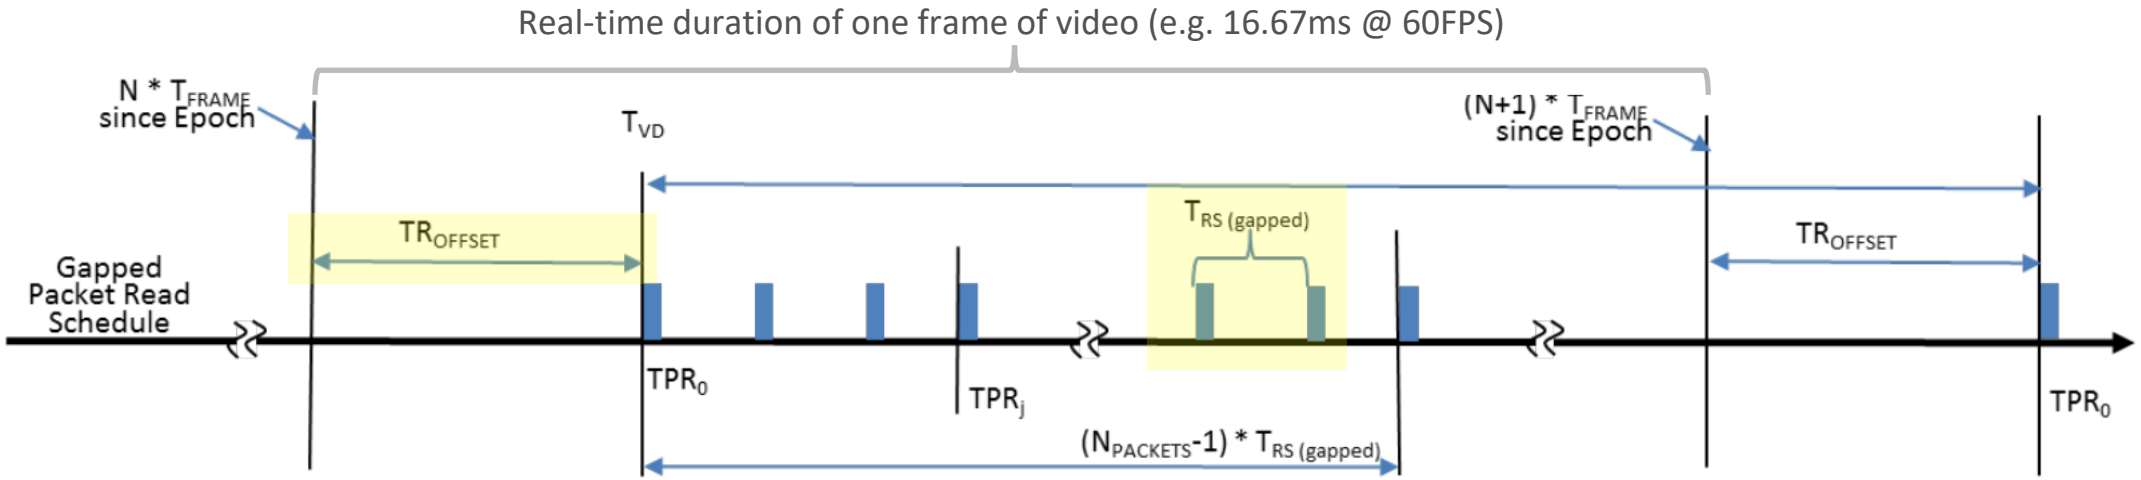

All images from: SMPTE ST 2110-21:2017, with clarifying annotations by the authors

Network Interface Controllers (NICs): *Rate Limiting Function*

- Ethernet Network Interface Controllers (NICs) commonly provide packet rate limiting functions
- Rate limiting insufficient to meet ST2110-21 strict packet timing requirements
- RMTS introduces two new algorithms that, together with rate limiting, can achieve ST 2110-21 conformance

RMTS packet timing algorithms

- RMTS introduces two new algorithms to generate accurate T_{RS} and TR_{OFFSET}

1. For T_{RS}: a credit-based profiling and dynamic padding algorithm that significantly improves the existing hardware time precision and granularity, enabling precise packet pacing
2. For TR_{OFFSET}: a dynamic tuning algorithm which sets up a warm-up packet queue before transmission and dynamically adjusts the queue elements on demand to improve the timing accuracy of the 1st packet in each video frame

Accurate inter-packet gap: T_{RS}

- Credit-based profiling and dynamic padding algorithm in sender
- Control plane: profiles the rate limiter and set the rate limiter bandwidth (BW)
  - Credit scoring created for the packet queue based on elapsed time between packets
  - If credit > 0, then NIC transmits packet and reduces credit
  - If credit <= 0, then queue is blocked
- Packets now go out at the rate limit ON AVERAGE
- Data plane: establishes the transmission packet pattern for additional timing precision
  - Pause packets (not transmitted) are added to the queue to create accurate packet spacing, so rate limit bursts are avoided

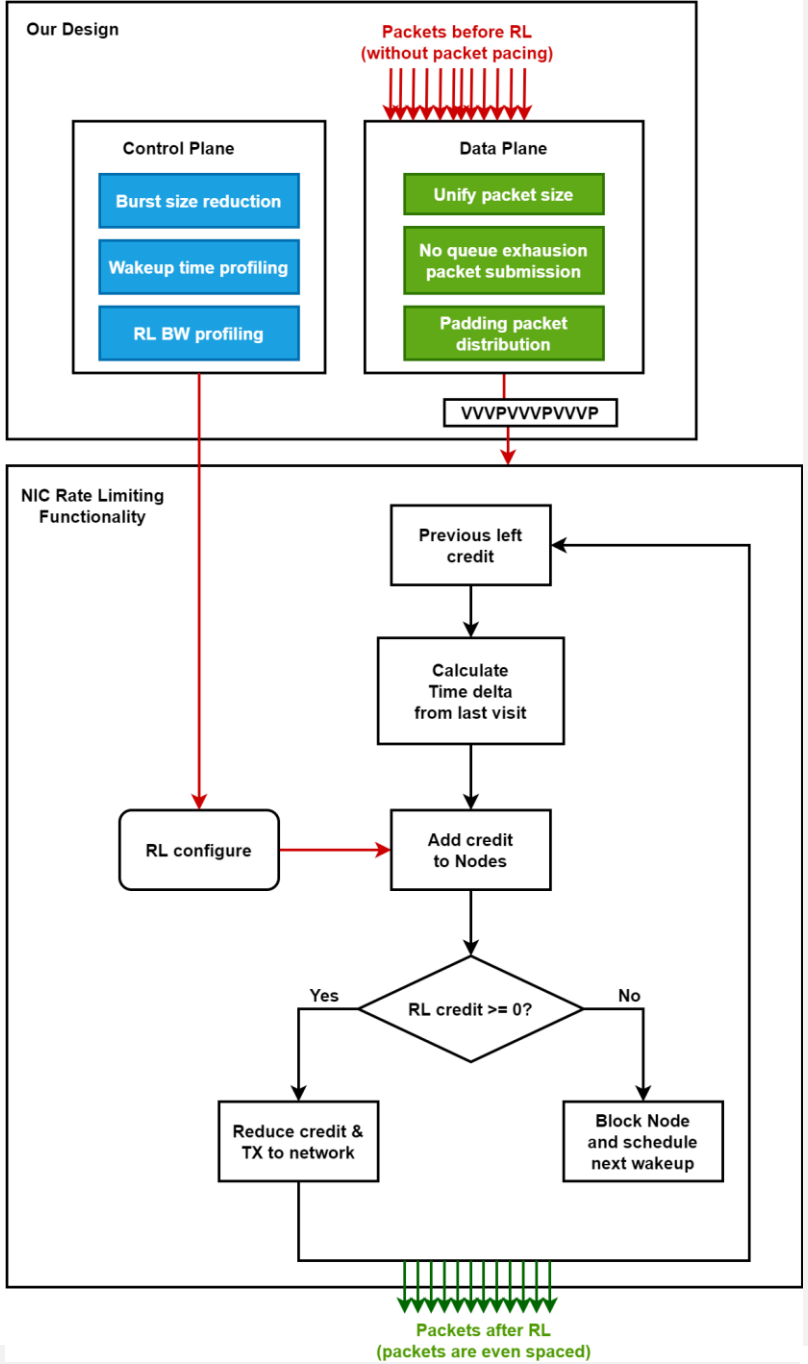

Accurate 1st packet timing: TR_{OFFSET}

- Launch time granularity limit
  - CPU sends timestamped packet to NIC
  - Packet buffered in NIC queue
  - NIC schedules packet for delivery (too late/too early)
- Solution
  1. Setup warm-up packet queue
  2. Collect actual packet transmit time
  3. Adjust warm-up queue packets until correct TR_{OFFSET} is achieved
- No longer counting solely on launch time accuracy

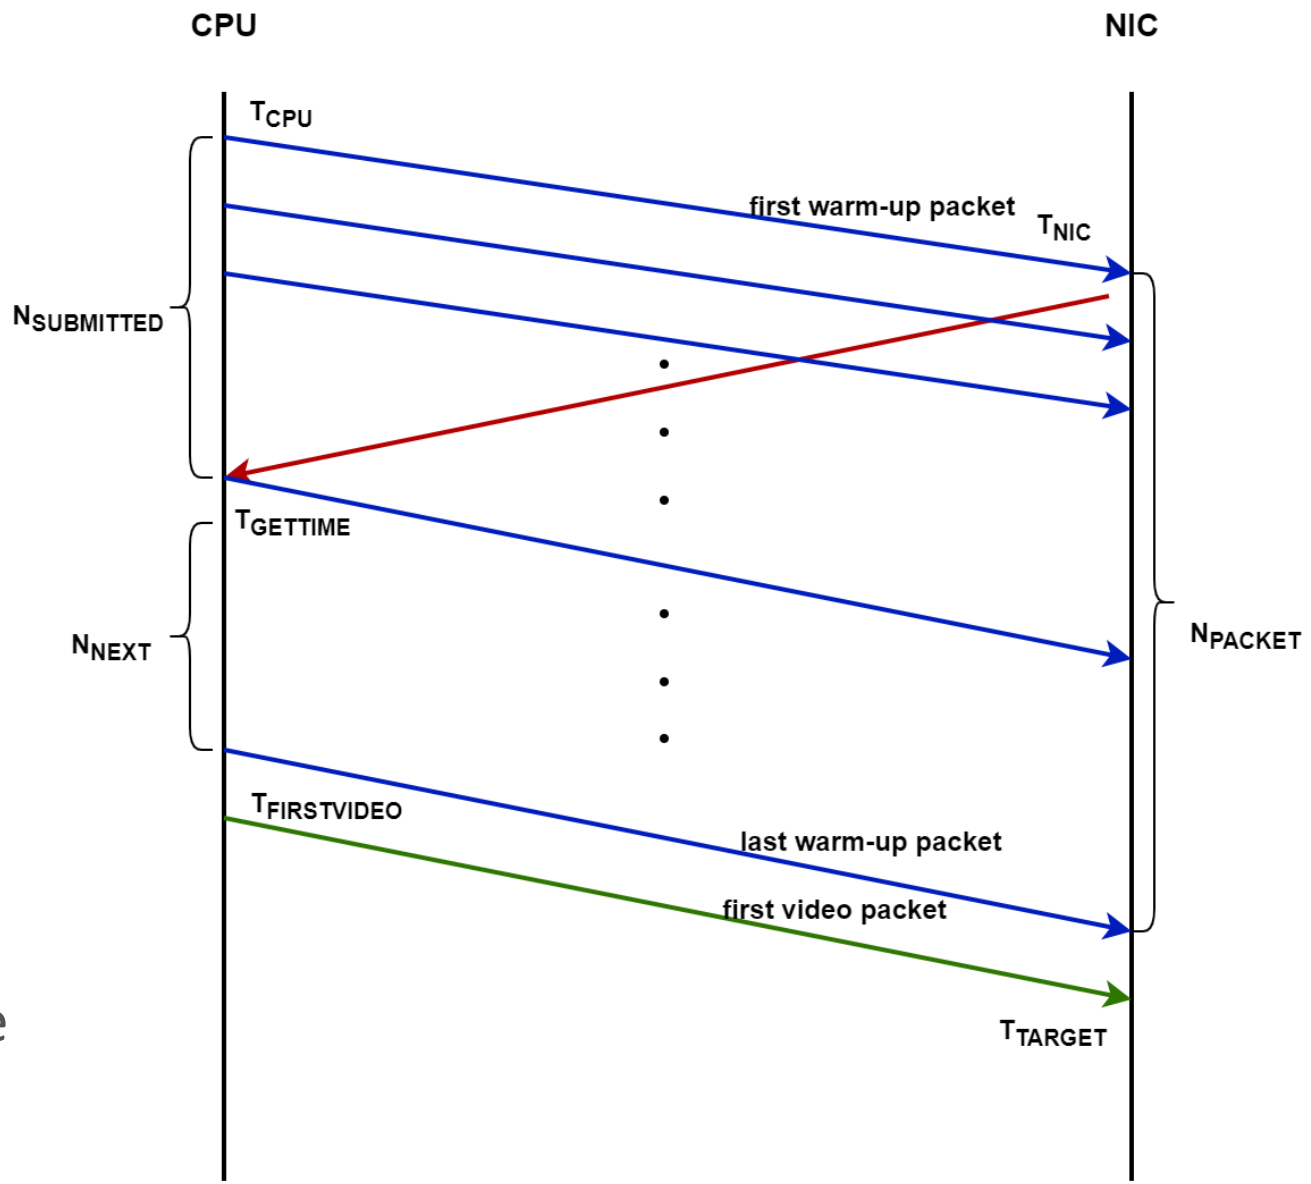

Packet pacing (T_{RS}): experimental results

- VRX_{FULL} is an ST2110-21 parameter (sender leaky bucket full) that is required to not exceed a maximum value
- From the test results:

  Fig 1: Rate limit set, no packet padding

  - Timing deviation accumulates exceeding VRX threshold

  Fig 2: Rate limit set with packet padding enabled

  - VRX stays below required threshold

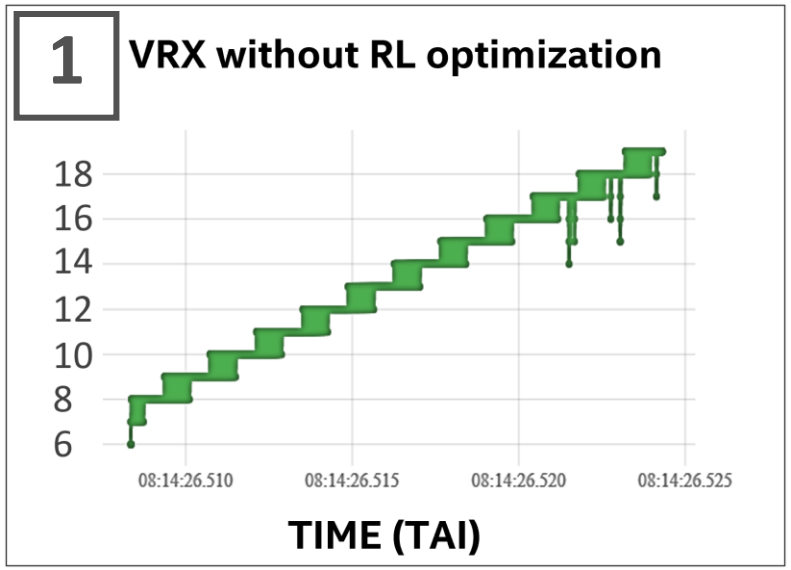

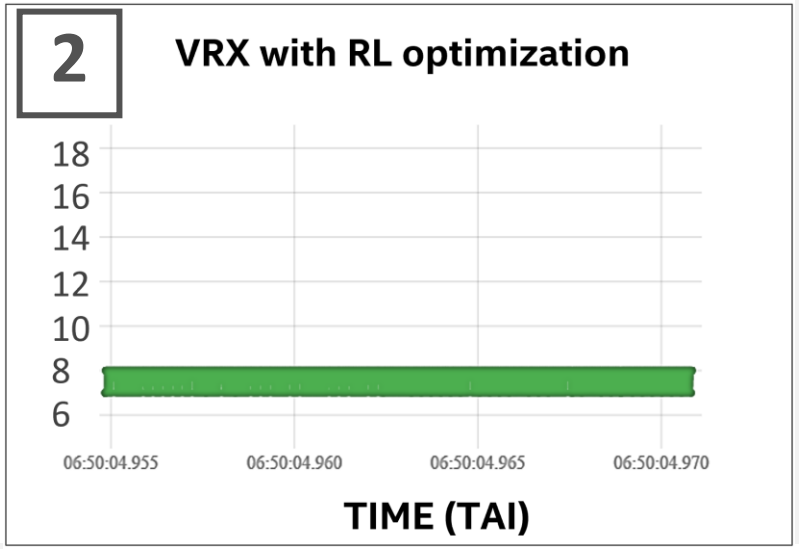

1st packet timing (TR_{OFFSET}): Experimental results

- From the test results:

Fig 1: Warm up queue disabled

- TR_{OFFSET} generally stays within the required range, but sometimes exceeds the timing by 30-40μs, failing ST 2110-21 compliance

Fig 2: Warm up queue enabled

- TR_{OFFSET} stays reliably within a narrow range, conforming to the ST 2110-21 standard

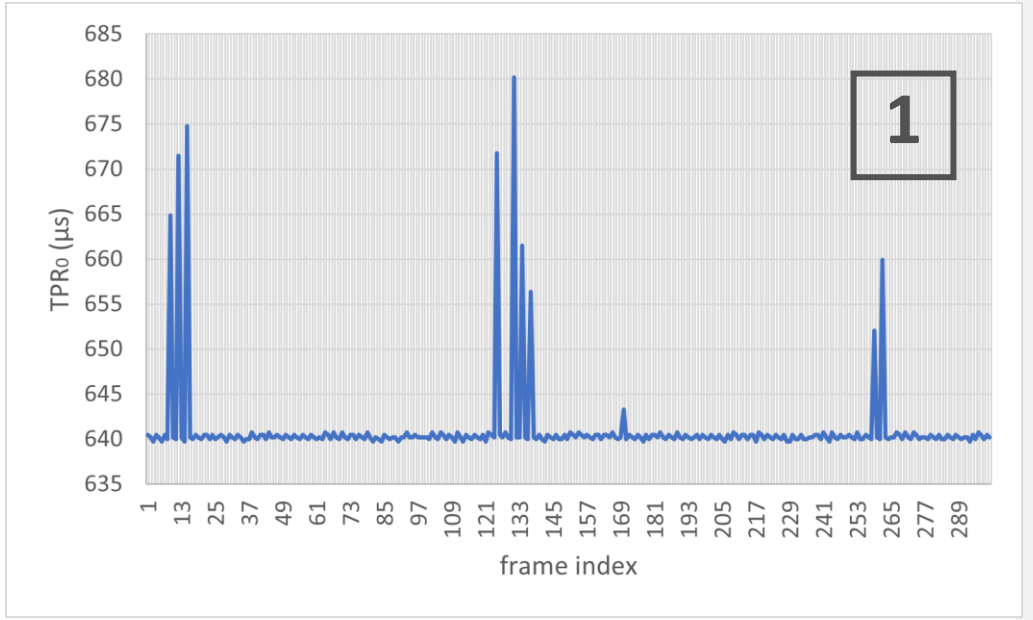

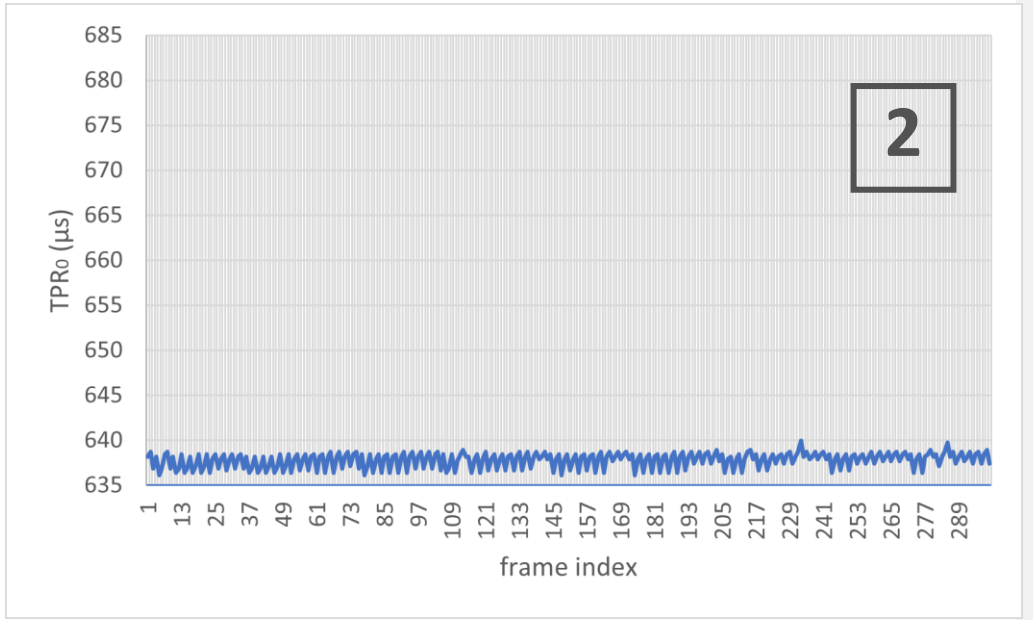

Conclusion

- RMTS has been tested with many resolutions, frame rates and video formats up to 8K, consistently passing all ST 2110 compliance tests
- Extensive ST 2110 equipment inter-operability tests also confirm these results
- RMTS offers a complete ST 2110 compliant implementation
- RMTS source code is available with a BSD-3 license at: https://github.com/OpenVisualCloud/Media-Transport-Library

THANK YOU