A Performance Measurement Study of the Reliable Internet Stream Transport Protocol

Ciro A. Noronha, Ph.D.
Director of Technology, Compression Systems
Cobalt Digital

Agenda

- Motivation
- Overview of the Reliable Internet Stream Transport protocol
- Performance Measurement
  - Packet loss performance
  - Packet re-ordering configurations
- Conclusions: how to fine-tune a RIST link
- Review of multi-company demonstrations

Motivation

- Advances in compression technology and in network infrastructure have made it possible to use the Internet as a low-cost contribution link
- The Internet drops packets, and a recovery protocol is necessary as every packet loss is a glitch
- There are many proprietary solutions on the market that do not interoperate
- The Video Services Forum (VSF) formed the Reliable Internet Stream Transport (RIST) Activity Group in early 2017 to create a common specification for a protocol suite to solve this problem
- RIST Simple Profile was published October 2018

Packet Recovery using ARQ

- ARQ stands for:
  - **A**utomatic **R**epeat re**Q**uest
  - **A**utomatic **R**epeat **Q**uery
- This is the generic name for a number of retransmission strategies in the face of packet loss
  - Standard TCP uses a couple of ARQ variants
- In video transmission, the most useful variant is "Selective Retransmission" (NACK-based)
  - If you don't hear from me, everything is OK
  - If I miss anything, I let you know and you resend just that
- RIST uses ARQ

ARQ Illustration

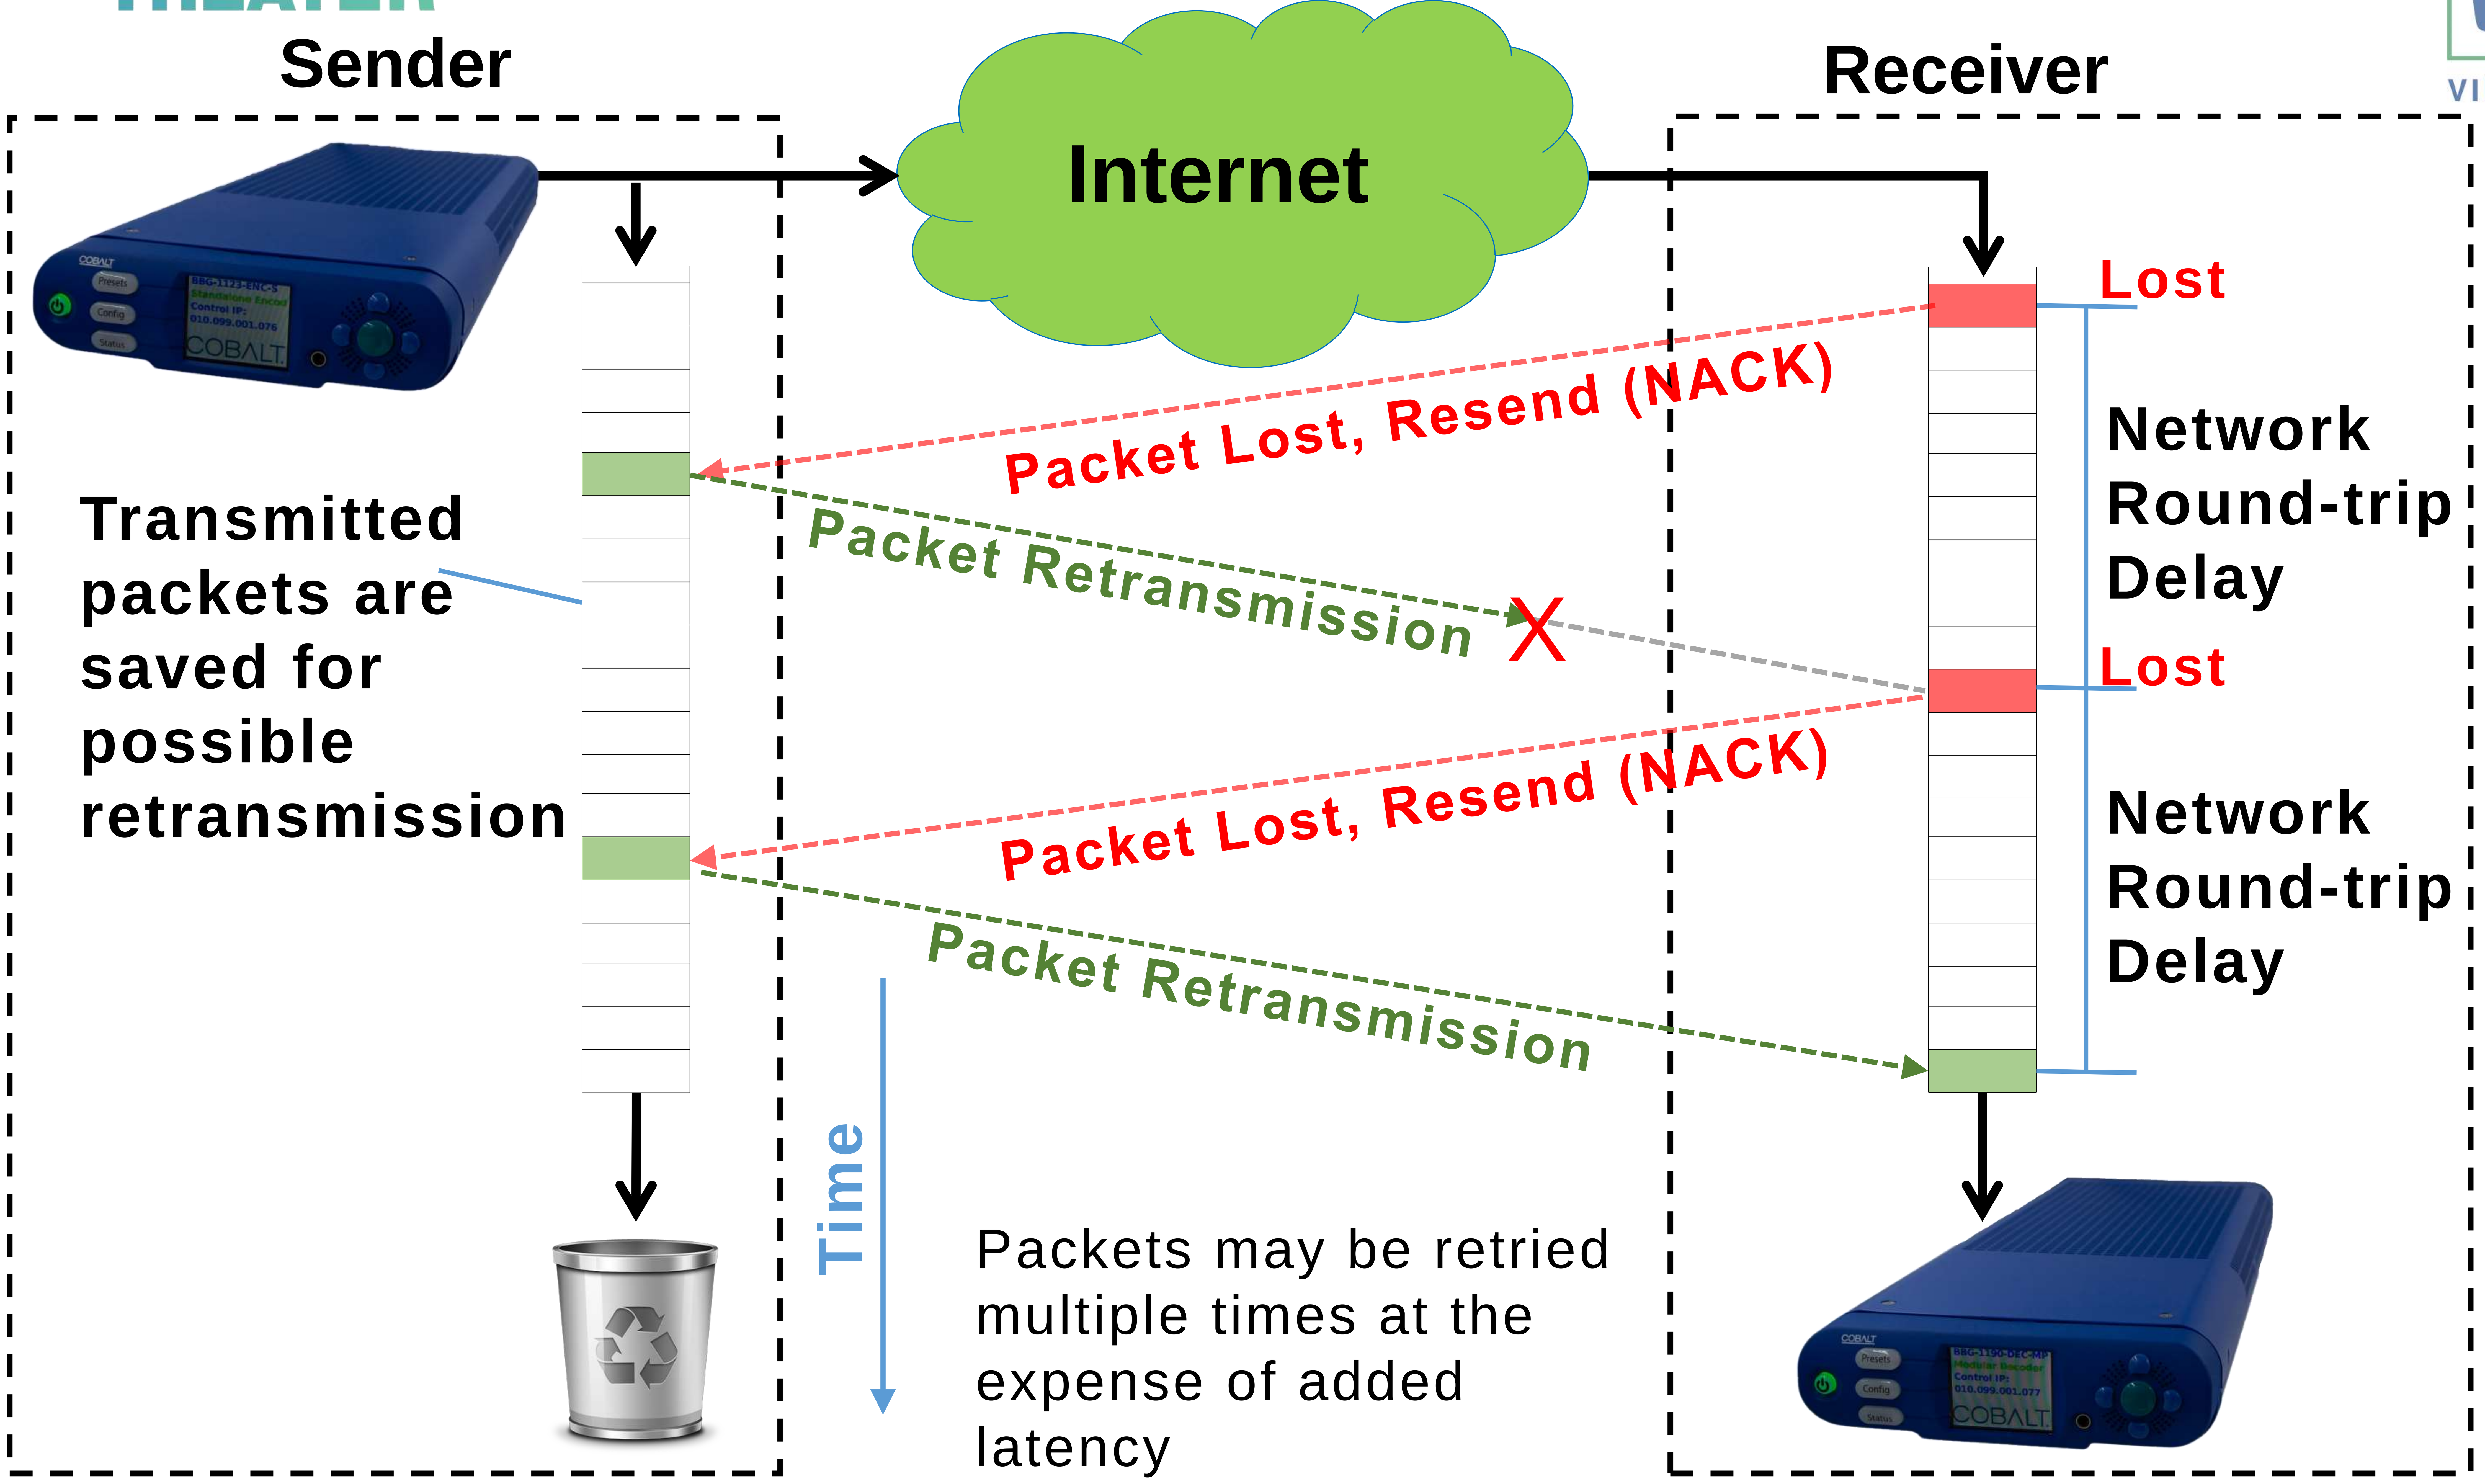

RIST Protocol Basics (Sender)

- Primary stream transmission is through RTP, using the relevant standards
  - SMPTE-2022-1 for Transport Streams
  - UDP flow sent to port P, where P is an even number
- RIST sender is required to transmit RTCP packets
  - Packets sent to port P+1
  - Primary function is to establish state in firewalls for the NACK return packets
  - Suggested content:
    - Sender Report (SR) plus CNAME
    - Empty Receiver Report (RR) plus CNAME

RIST Protocol Basics (Receiver)

- Receiver listens on port P for the content, and on port P+1 for the RTCP packets
- Receiver sends periodic RTCP packets (RR+CNAME)
  - Receiver RTCP packets are sent to the source IP address and source UDP port of the received RTCP packets
  - Firewalls will treat these as “response” to the sender RTCP packets
- If the receiver detects packet loss, it will send a retransmission request for the missing packets
  - Retransmission request is an RTCP packet

RIST Retransmission Requests

- RIST NACKs (Retransmission Requests) are built using standard compound RTCP packets
- A compound RTCP packet from a RIST receiver will contain RR (may be empty), CNAME, and NACK.
- RIST has defined two types of NACK messages:
  - Bitmask Message:
    - Can request any pattern within a group of 17 consecutive packets
    - Useful for "salt and pepper" loss
    - Generic NACK from RFC 4585
  - Range Message
    - Can request a block of consecutive packets
    - Implemented with Application-Defined RTCP message
    - RIST AG may approach IANA for a permanent registration

RIST Retransmissions

- RIST retransmissions are an exact copy of the original missed packet
- Retransmitted packets are sent together with media packets (RTP sent to the same port P)
- Retransmitted packets are differentiated from original packets using the SSRC field
  - Last bit of SSRC is **zero** for original packets, **one** for retransmissions
  - Identifying retransmissions helps with system stability

RIST and Firewalls

Bonding Support

- RIST Simple Profile has support for Bonding
  - Sender splits the stream over multiple physical channels
  - Receiver can send NACKs over each of the paths
  - Can also be use for redundancy (in the same fashion as SMPTE-2022-7)
    - Two or more copies of the same stream can be sent over distinct links

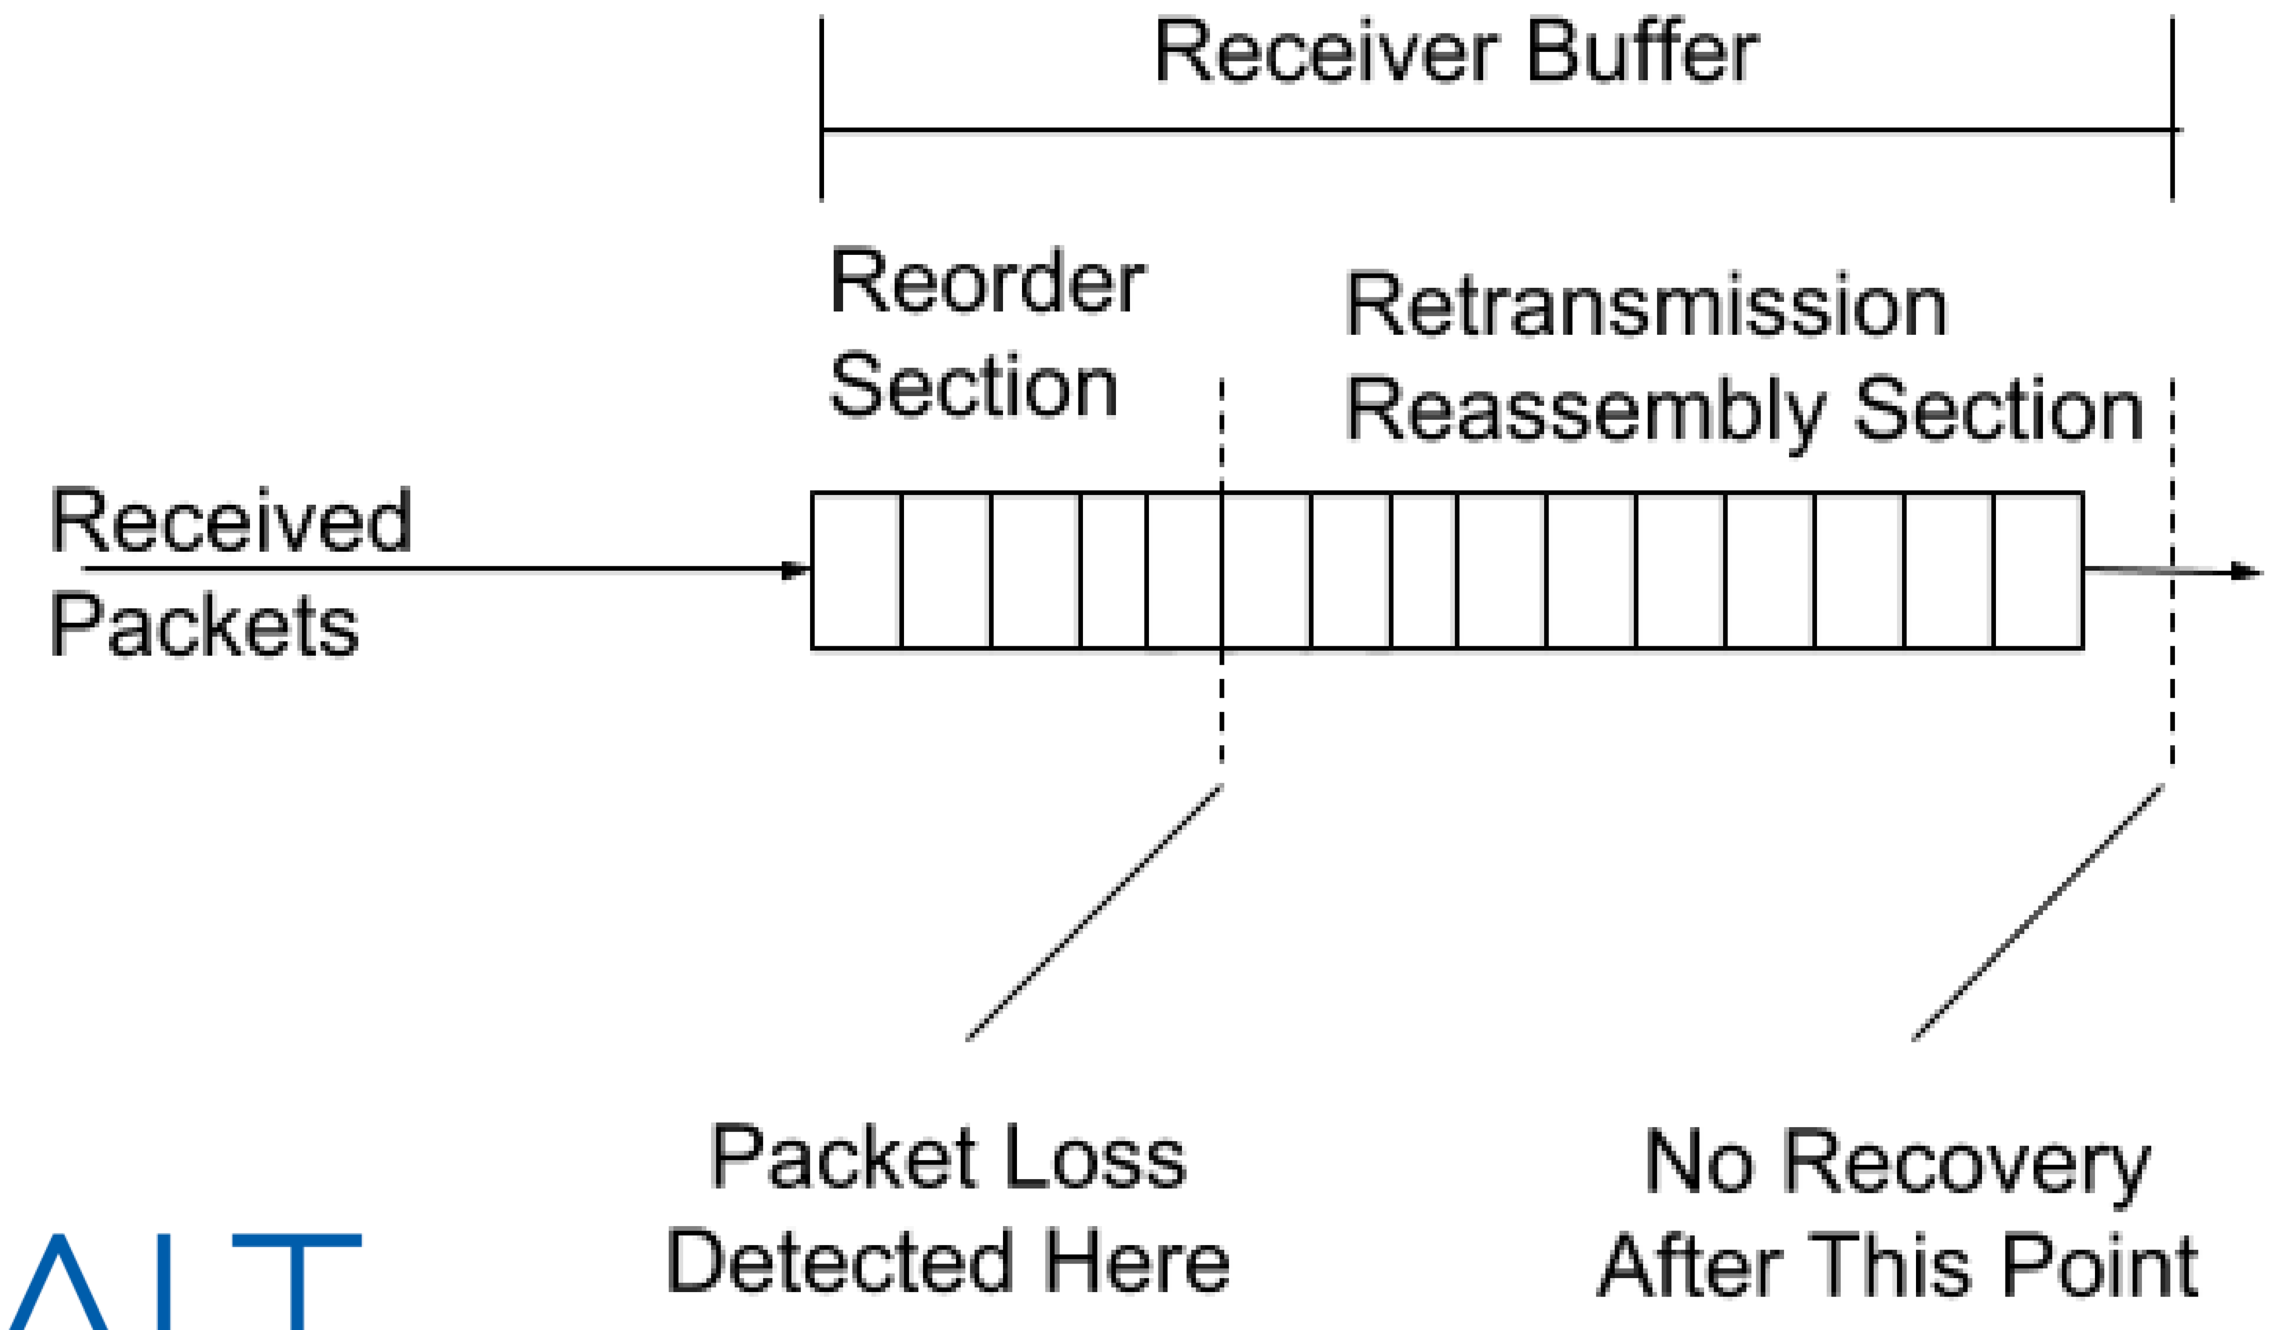

Packet reordering is supported by adding a reorder section to the receiver buffer

Packet Loss Performance Measurement

Network Emulator

Encoder

Decoder

Packet Delay
Linux “netem”

Packet Dropping
Custom App

Test Automation
Custom SNMP

- Media bit rate: 8 Mb/s (1920×1080i59.54 source)
- Simulated round-trip delay: 200 milliseconds
- Random i.i.d. packet losses:
  - Single packet losses
  - 5-packet burst losses
- Two-minute runs
- Independent variable: number of retries, tested from 1 to 10
- Receiver retransmission buffer set to (200*R* + 100) milliseconds, where *R* is the number of retries
- Sender buffer set high enough to handle the worst-case receiver buffer
- For each retry value, increase the packet loss until at least one unrecovered packet is detected in the two-minute run.
- Record this packet loss rate
- Repeat each test 10 times

Single-Loss Results

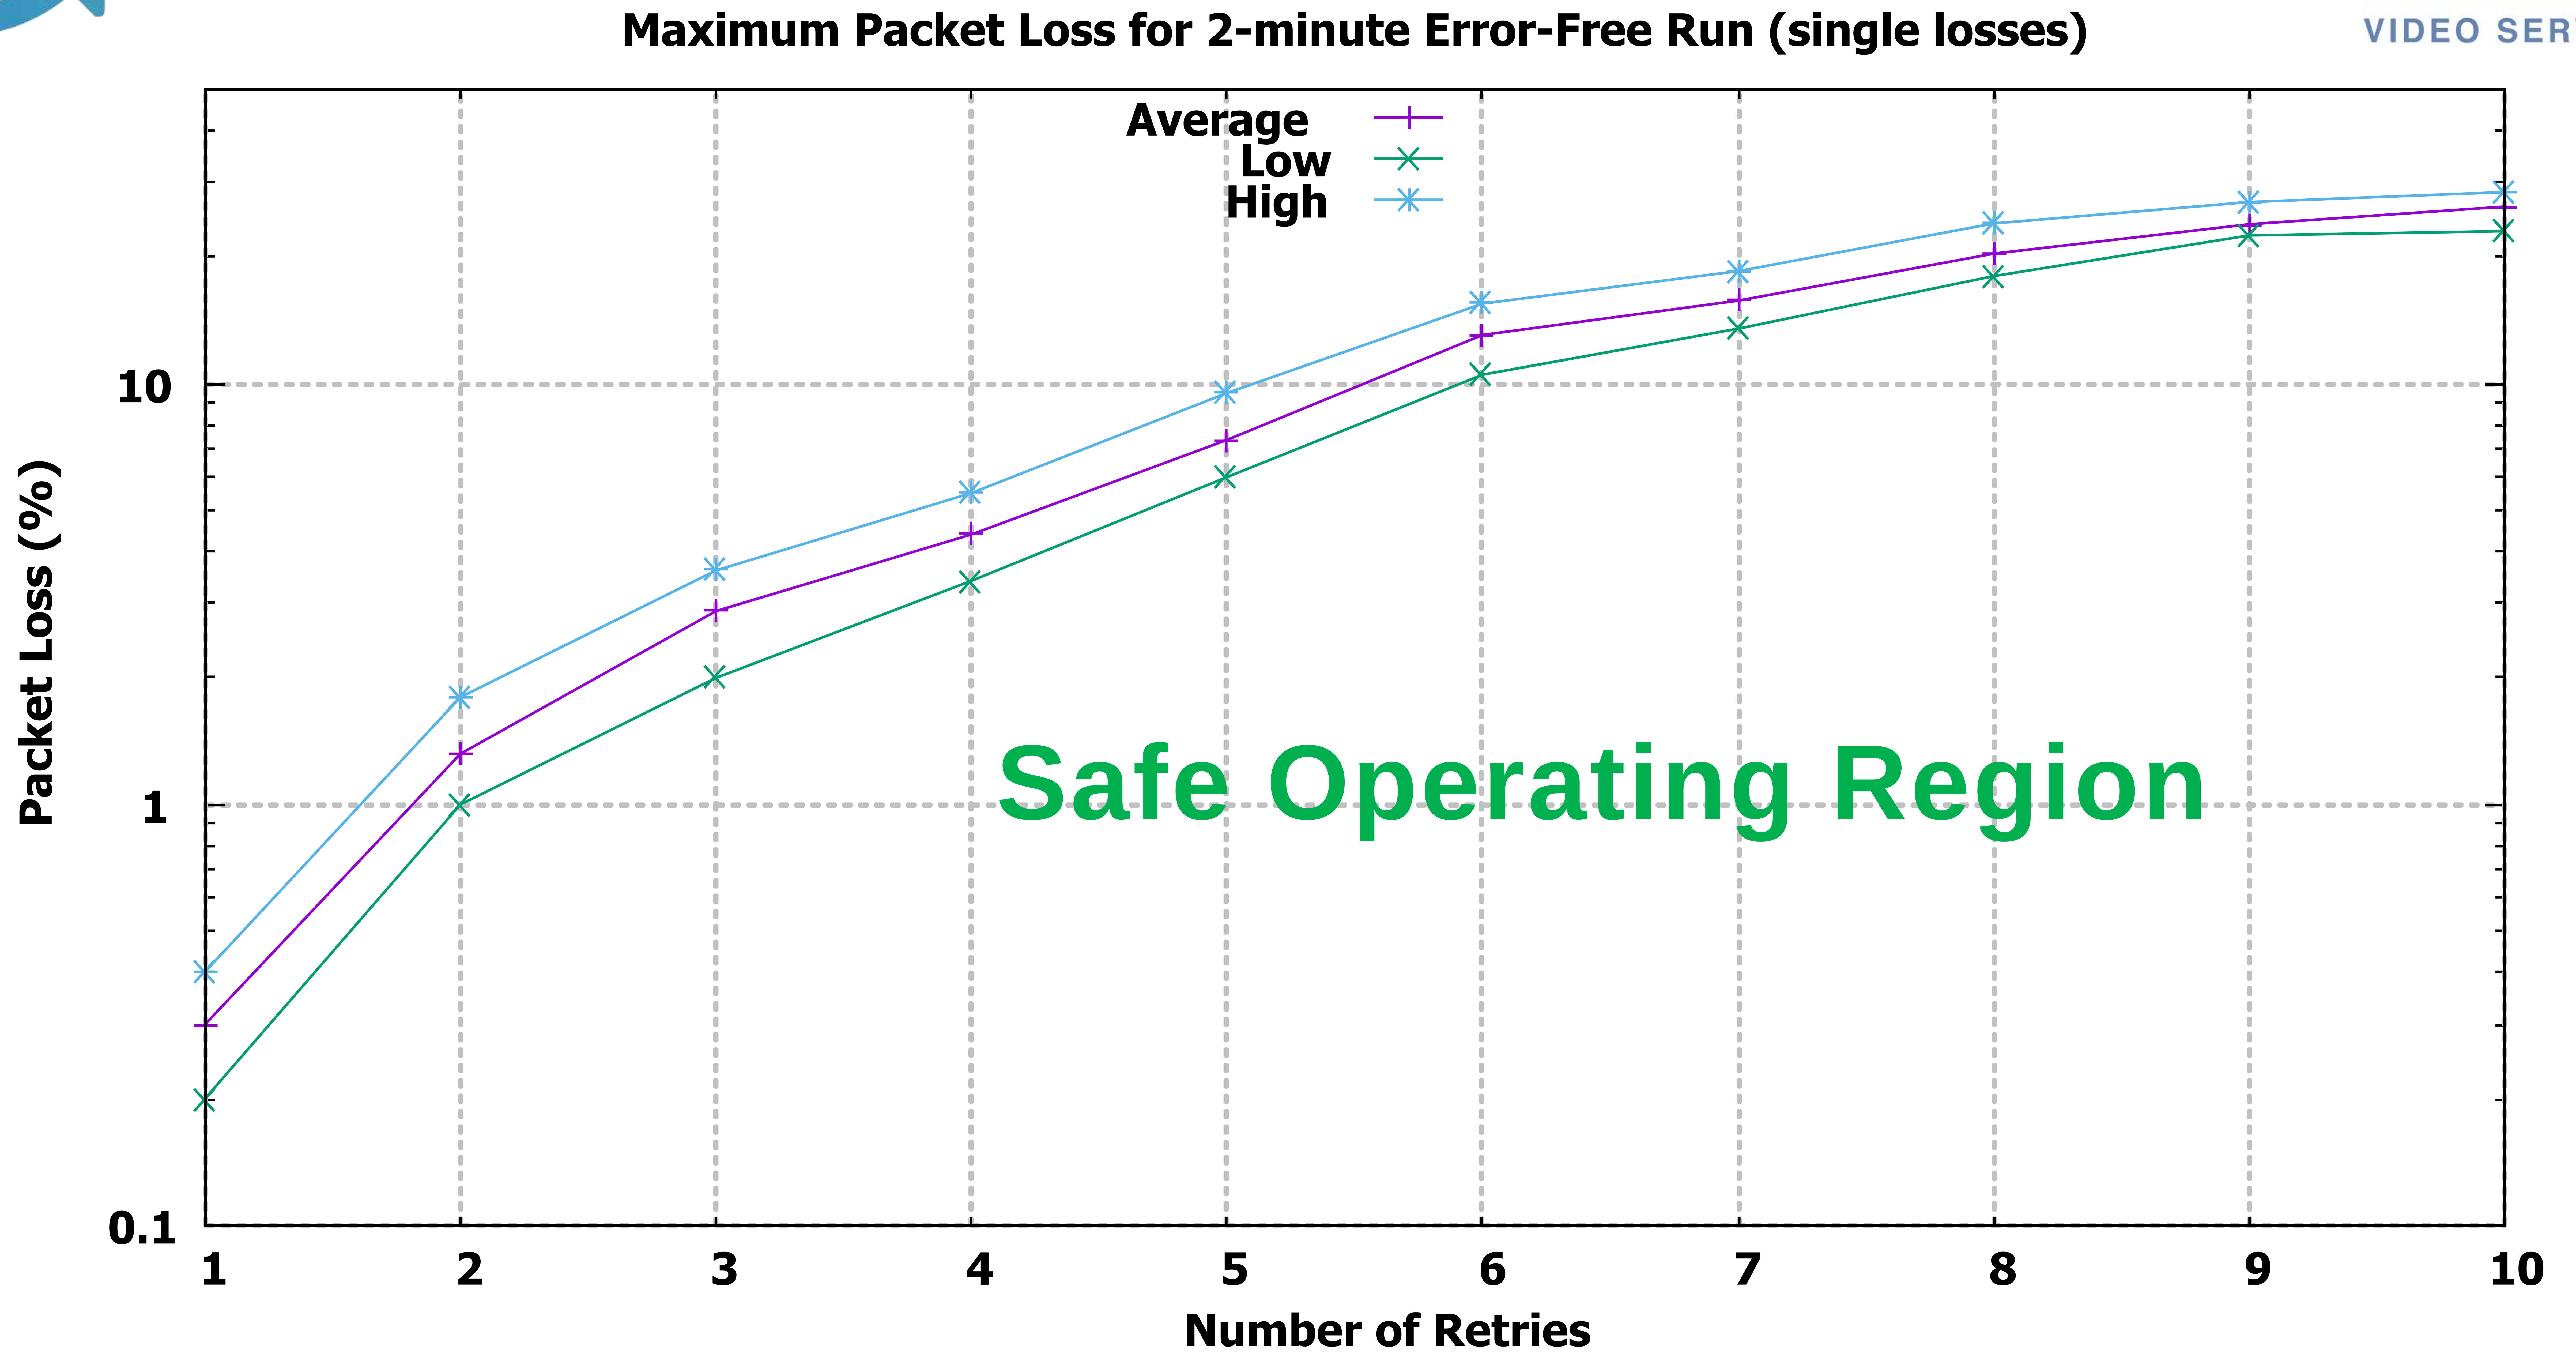

Burst Loss Results

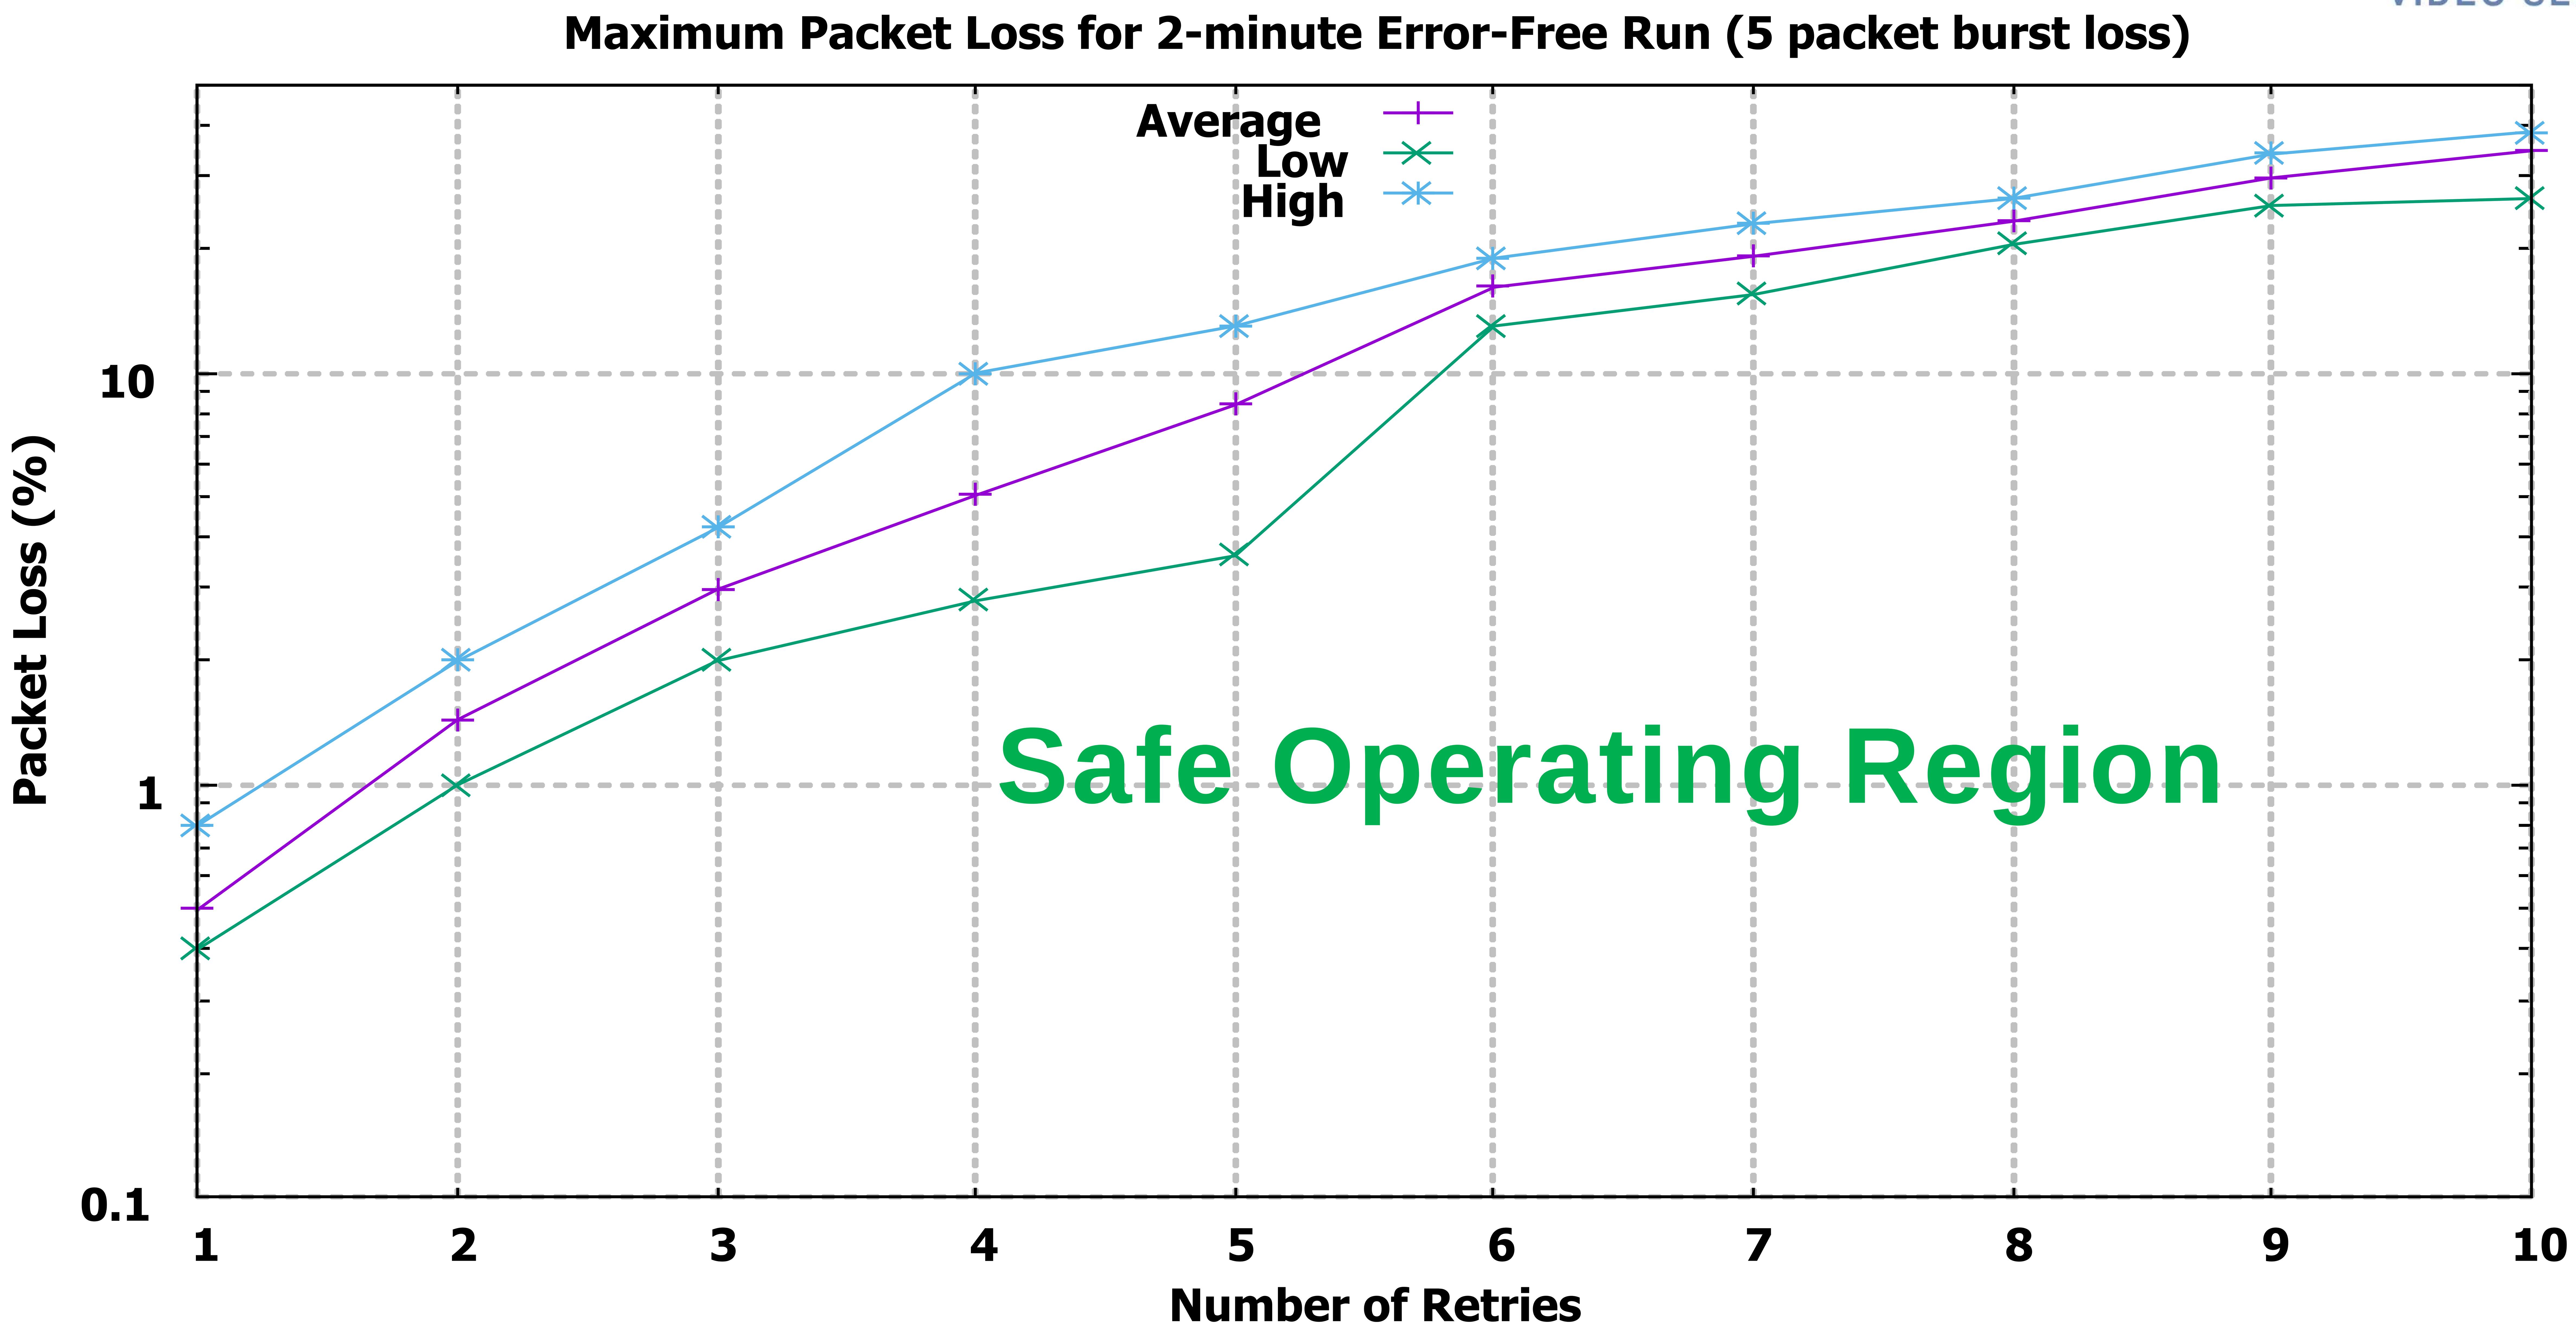

Packet Re-Ordering

- In the Internet, packet re-ordering only happens when paths change
  - The only way a packet with "overtake and pass" another is if it uses a different (shorter) path
- Question: if not using bonding or multipath intentionally, is it necessary to accommodate packet re-order?
- Trade-offs:
  - Non-zero re-order buffer: increased latency
  - Zero re-order buffer: possibility of unnecessary retransmissions
- Question can only be answered with actual data on Internet traffic

Data from the Internet

	Total Packets	Reorderings	% Reorder
CDN	90,905,926	28,558	0.031%
Tier-1 ISP	39,403,671	307,615	0.781%
Tier-2 ISP	245,535,161	943,188	0.384%
OC48	153,143,822	653,717	0.427%
Total	**528,988,580**	**1,933,078**	**0.365%**

- Internet backbone measurements indicate that the incidence of out-of-order packets is, on average, a fraction of a percent of the traffic.
- In the absence of any additional information, it is unnecessary to set a re-order buffer for a single-link RIST connection over the Internet.

Data derived from:

Jaiswal, S., Iannaccone, G., Diot, C., Kurose, J., and Towsley, D., "Measurement and Classification of Out-of-Sequence Packets in a Tier-1 IP Backbone", *IEEE INFOCOM 2003*, San Francisco, April 2003.

Configuring a RIST Link

- Input parameters/requirements (site data):
  - Network round-trip time (found with "`ping`")
  - Maximum acceptable transport latency (if required)
  - Network loss (if known)
- Configurable parameters:
  - Retransmission Buffer
  - Re-order Buffer
  - Number of Retries
- Problem: select the values for the configurable parameters from the site data

Recommendations

- If there is a latency limit:
  - Set the retransmission buffer to the latency limit
  - Divide the latency limit by the round trip time and round up to find the number of retries
- If there is no latency limit:
  - If the network loss is known, read the number of retries from the performance plots and add a margin; set the retransmission buffer to at least the number of retries times the round trip
  - If the network loss is not known, a good starting point for the number of retries is 4
- Set transmitter buffer size (if configurable) as high as it will go
- Re-order can be set to zero unless using bonding
  - If using bonding, set to at least the worst case differential delay

IBC 2018 Demo

- 8 companies each sent a 5 Mb/s stream over the Internet to the Cobalt headquarters in Champaign, Illinois
- The streams were received by Cobalt 9990-DEC decoders, combined in a multiviewer, and published to YouTube in real time
- Streams were sent from UK, Canada, Israel, and the US (Northern CA, Southern CA, Florida, Virginia and Massachusetts)
- Independent implementations from the specification (no source code sharing)


Israel
VideoFlow
USA - CA
Cobalt
USA - CA
QVidium
United Kingdom
BUILDING THE CLOUD OF
REAL-TIME
Nevion
RIST
RELIABLE INTERNET STREAM TRANSPORT
LIVE
Interop Demo
For IBC 2018
Receivers and Live Composite Stream
Provided by:
COBALT
VSF
VIDEO SERVICES FORUM
Receivers: 9990-DEC-MPEG
Location: Champaign, Illinois, USA
USA - VA
Zixi
USA - FL
DVEO
Canada
Evertz
USA - MA
Broadcasting/CATV
Artel

VidTrans 2019 Demo

- VidTrans 2019 was held in Los Angeles (Marina Del Rey) in February 2019
- A number of participating companies provided on-site receivers at the conference
- Streams were sent from locations in the world to the receivers at the conference
  - "Mix and match" of senders and receivers
- A camera in the show floor transmitted to a relay in the San Francisco area which bounced it back to the conference
  - Sub 1-second end-to-end latency


Sweden
Net Insight to VideoFlow
Low Delay Loop to Northern CA
Cobalt to VideoFlow
Israel
VideoFlow to Cobalt
USA - Massachusetts
Government
Artel to Cobalt
USA - San Diego
QVidium to Nevion
Israel/France
VideoFlow Bonding
USA - Florida
DVEO to Nevion
United Kingdom
SHAPING A NEW
FREEDOM
Nevion to Evertz
Canada
Evertz to Net Insight

Ongoing RIST Work

- Planned for future RIST profiles:
  - Content encryption
  - VPN support
  - NULL packet suppression (for transport streams)
  - Encoder rate control based on network availability
  - Support for high bit rate streams
- The objective is to provide all the features required for Internet contribution

Thank You

Ciro A. Noronha, Ph.D.
Cobalt Digital
ciro.noronha@cobaltdigital.com
+1 650 208-0605