Video Transport, Distribution and Quality of Experience

Ali C. Begen

Video and Content Platforms – Research & Advanced Development

Global IP Traffic Growth

IP traffic will increase 6x from 2007 to 2012

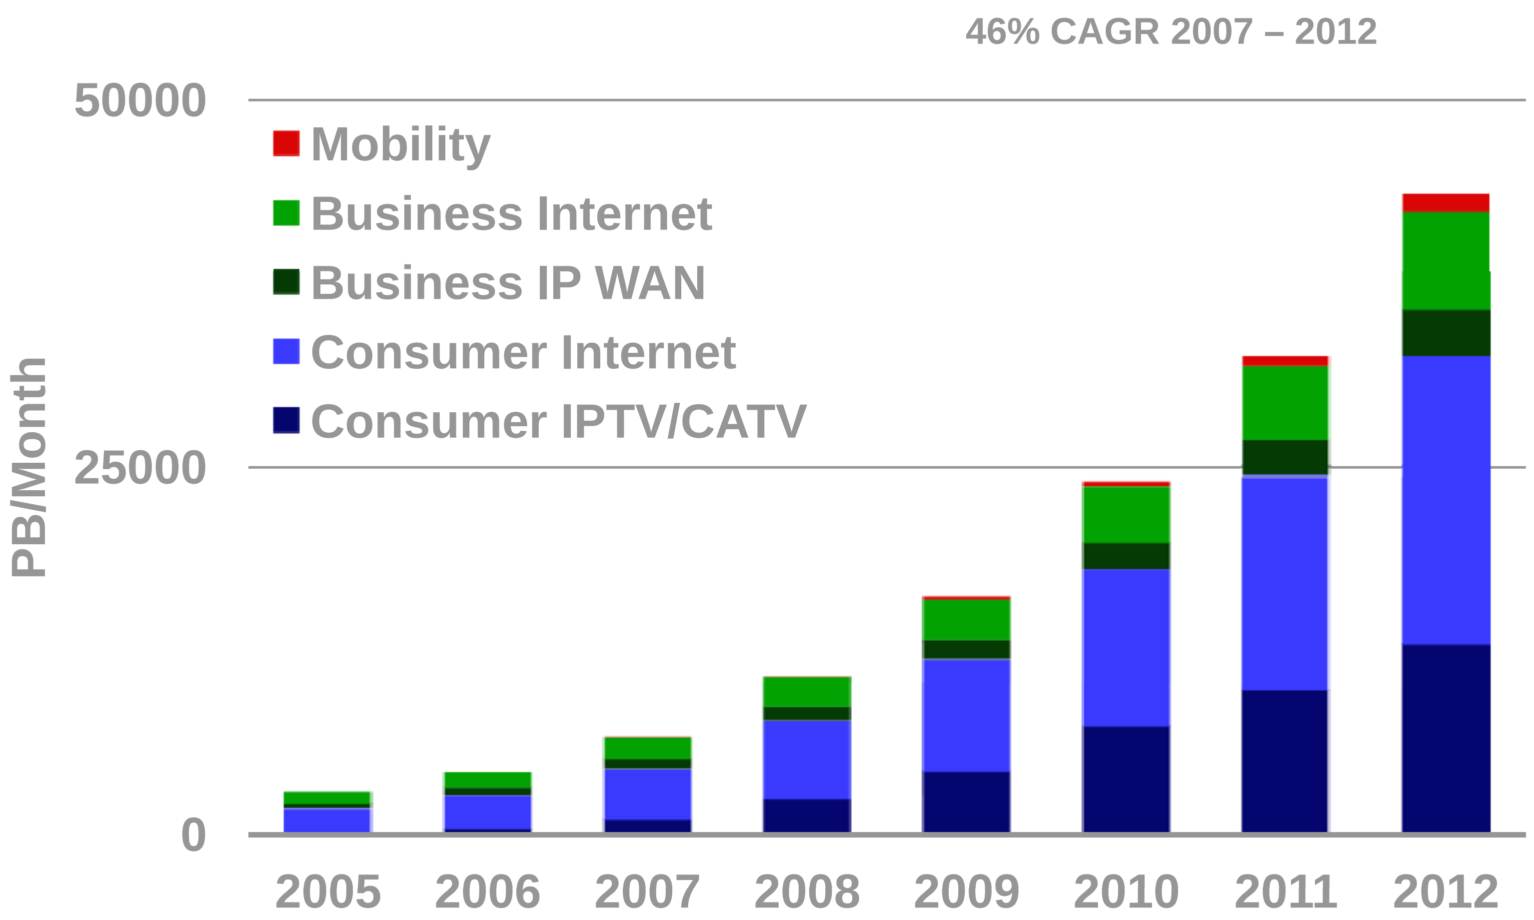

Petabyte: 1e15 bytes

http://www.cisco.com/go/ipngn

Household (US) Bandwidth Needs in 2010

1xHDTV + 1xSDTV + 2xPVRs + 1xVoIP + 2xPCs w/ HSD

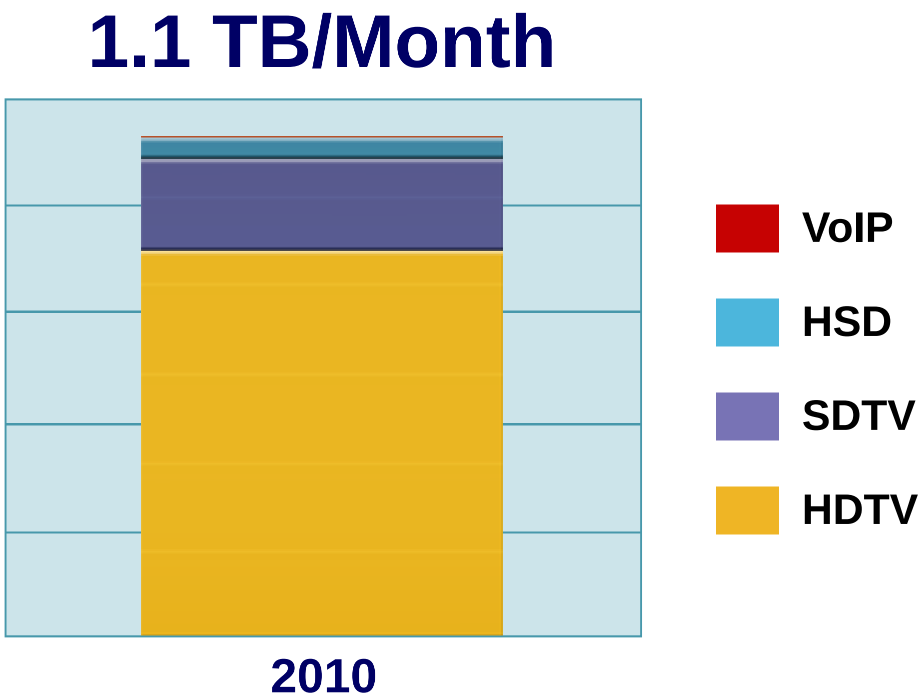

Twenty such homes would generate more traffic than traveled the entire Internet backbone in 1994/1995

YouTube and Hulu Traffic

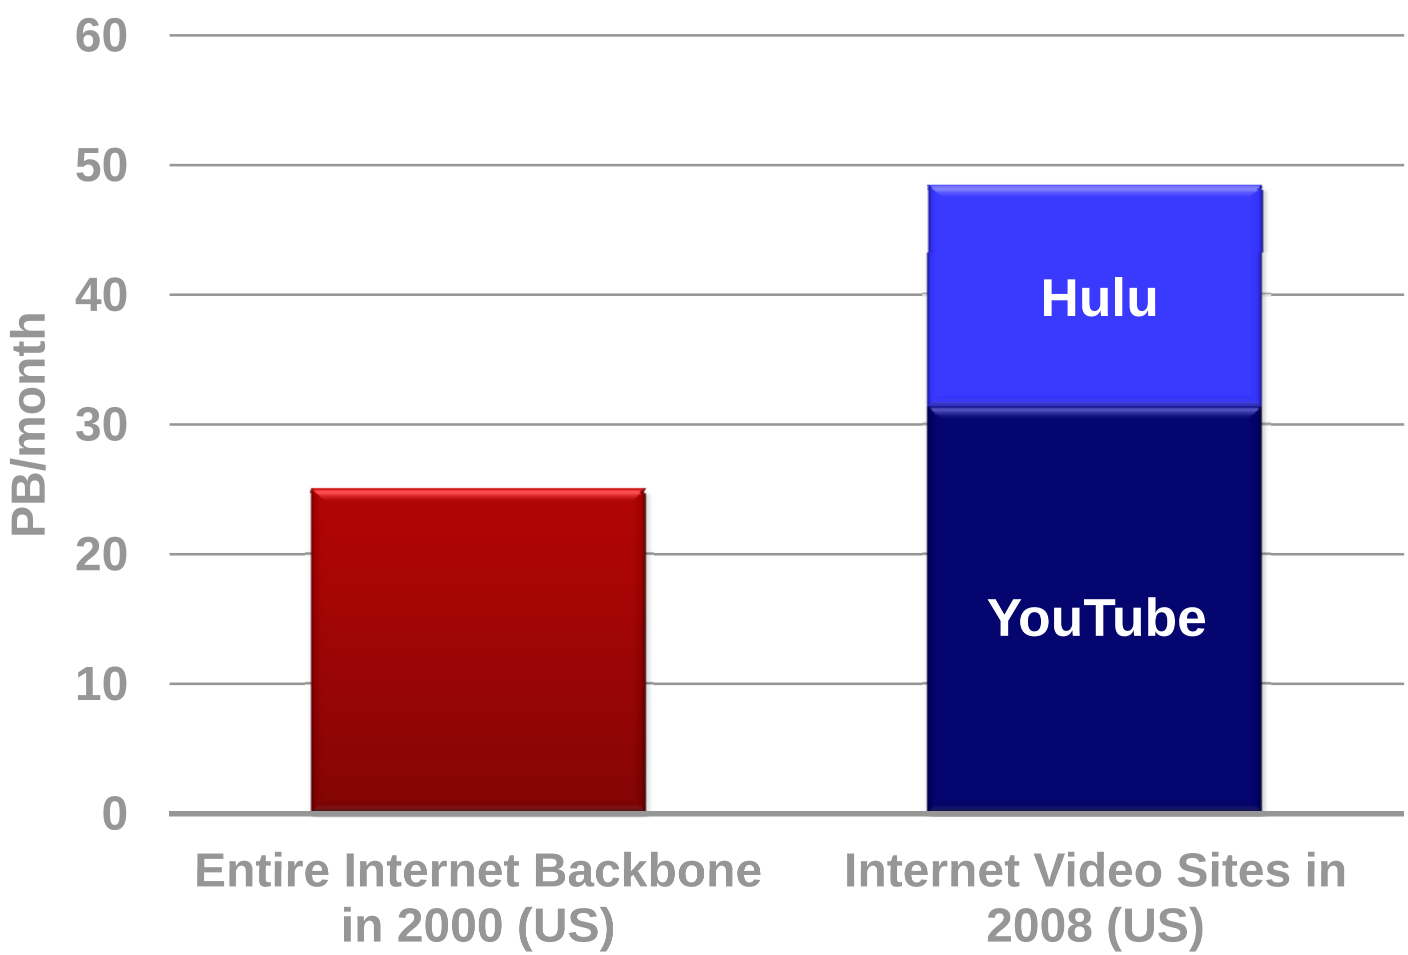

Petabyte: 1e15 bytes

Today's Outline

- IPTV architecture
    - Contribution
    - Distribution
    - Lossless IPTV transport
- Loss-repair methods
    - Forward Error Correction
    - Retransmission
- Channel changing in IPTV
    - The problem and its solution
    - Early results
    - Standardization efforts
- QoS/QoE monitoring

What Is IPTV?
The Fundamental Component for Connected Homes

IPTV = IP Network-Delivered TeleVision

- Switched digital video (SDV)
- Electronic program guides (EPG)
- Digital video recorder (DVR/PVR/nPVR)
- Video-on-demand (VoD)
- Interactive TV applications
- Targeted or advanced advertising

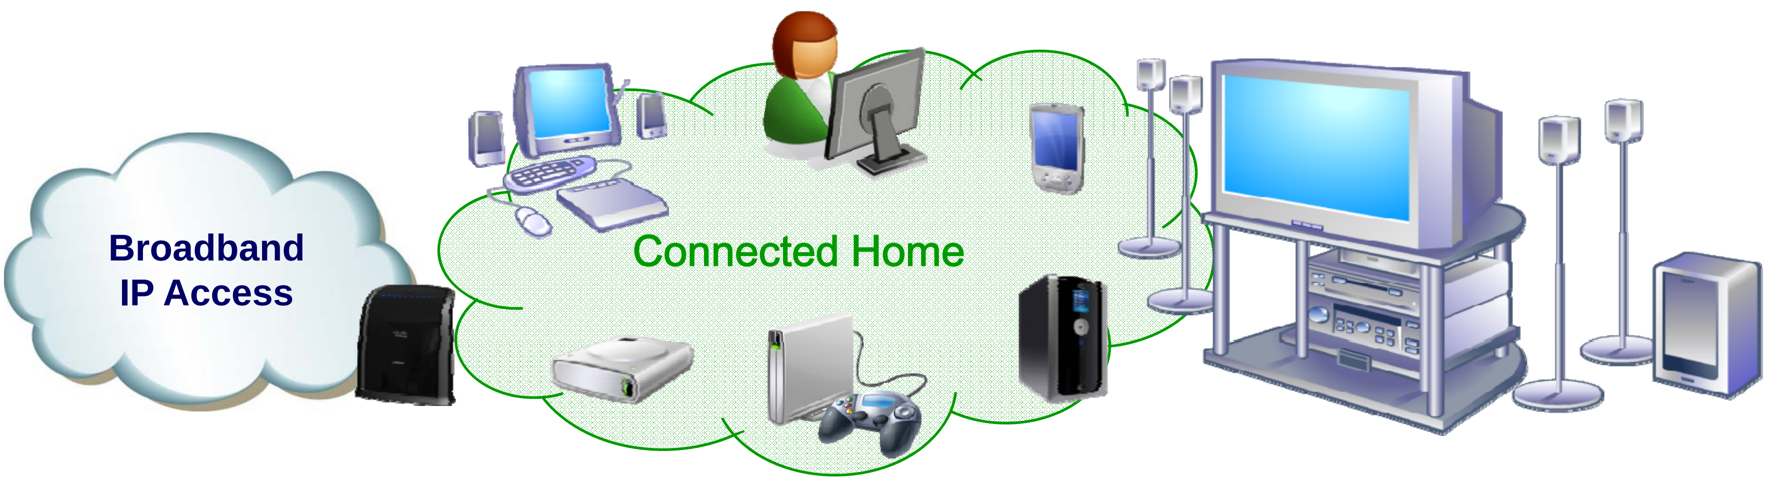

Trends Driving IPTV Adoption

- **Subscribers want more choice and control**
  - New generation grew up computer/Internet savvy
  - Connected Life – At home, at work, on the road
  - Want one bill, one provider, integrated services – Customized for me
- **Improved codec, access, server, & CPE technology**
  - MPEG-4 AVC (H.264) next generation codec improvements
  - New ADSL2+, VDSL2, FTTx, DOCSIS 3.0 access technologies
  - Moore's law advancements in processing & memory
- **Greater competition among service providers**
  - No longer limited by access – All services over any network
  - Traditional markets going away – Voice & long distance are almost free
- **Video is driving next generation SP network design**
  - Driven by video's bandwidth & QoS requirements
  - Experiencing exponential growth in Internet video usage

End-to-End IPTV Network Architecture

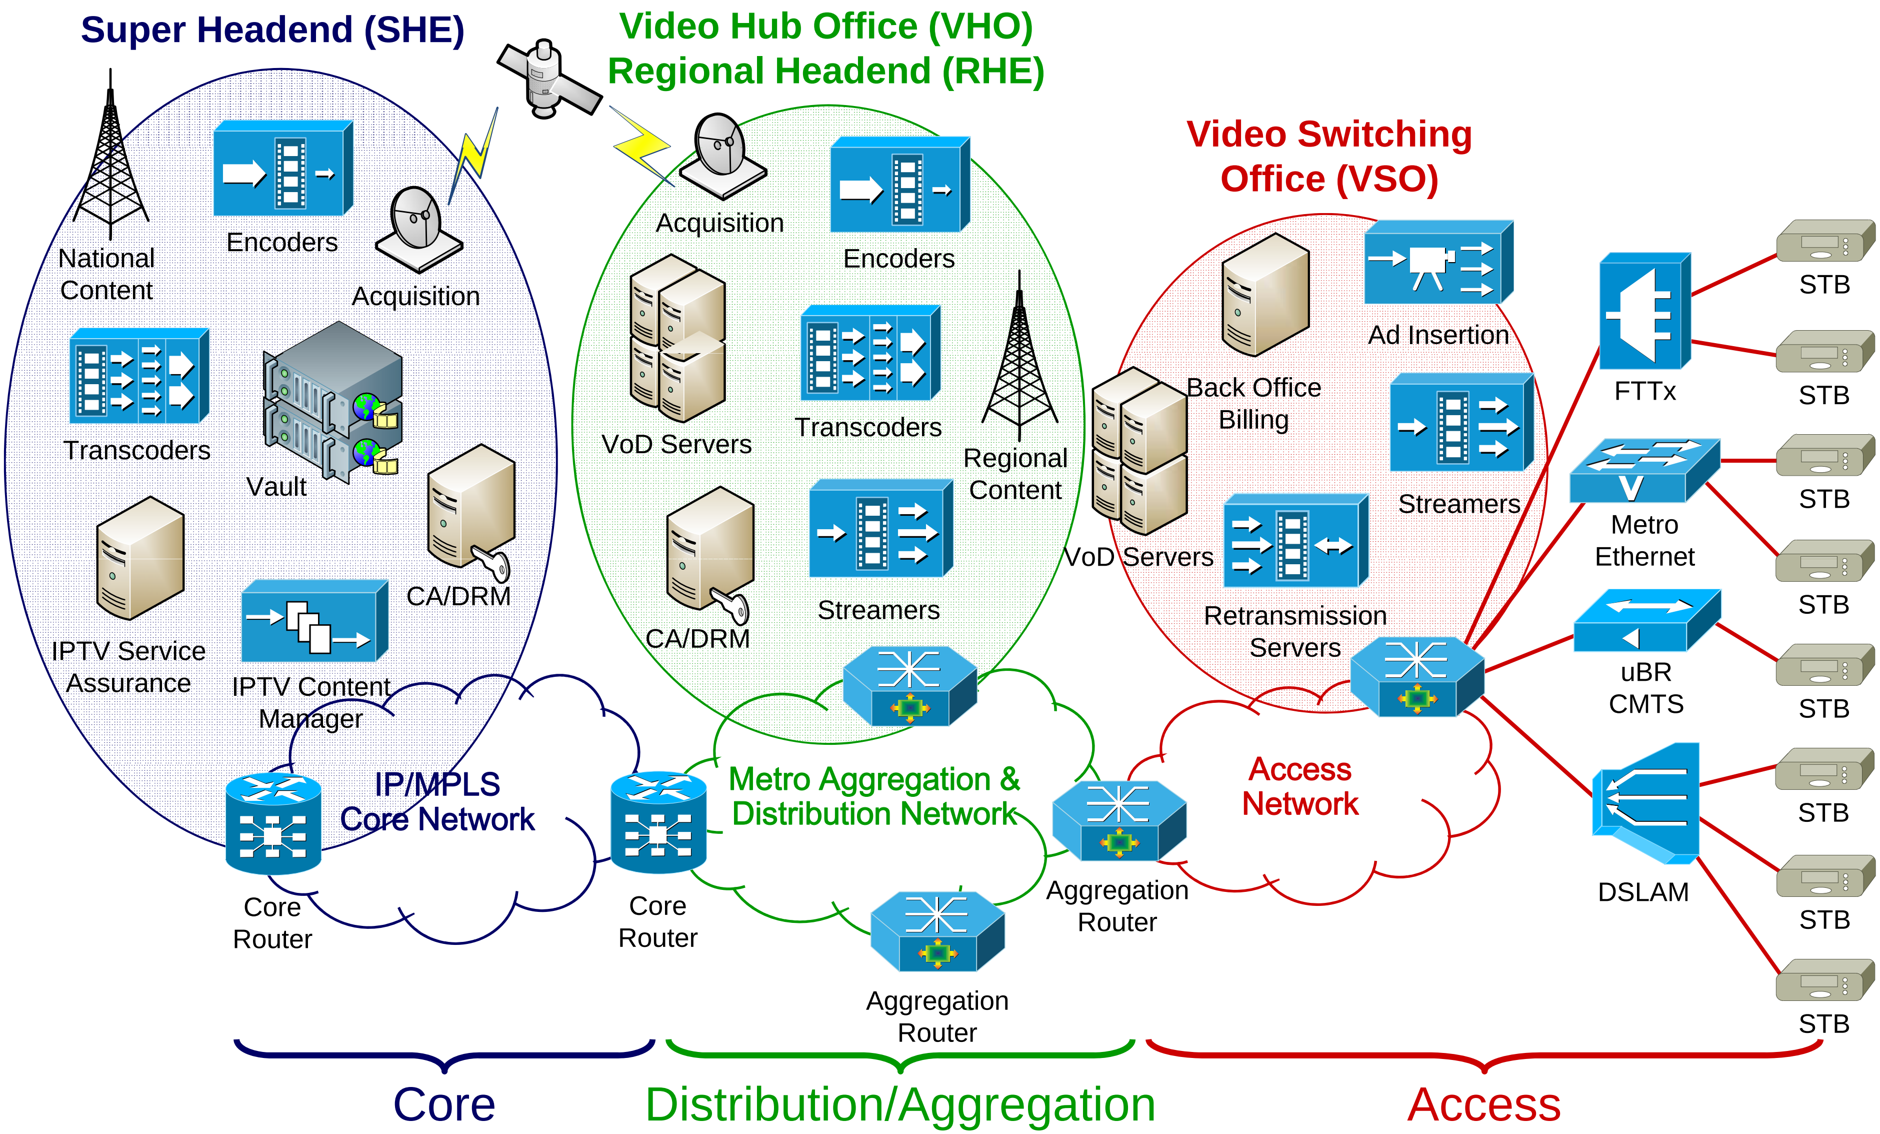

Types of Video Services

- **Transport (Contribution and Primary Distribution)**
- **IPTV (Secondary Distribution) / CATV**
  - IP multicast distribution from centralized super headends
  - Driving enhanced multicast features and functions
- **VoD (Secondary Distribution)**
  - Distributed architecture for better scalability
  - Non-real-time content distribution to caches
  - More impact on metro and access networks, less impact on the core
- **Enterprise**
  - mVPN based
  - Driving enhanced multicast features and functions
- **Over-the-Top (e.g., YouTube, AppleTV, Netflix)**
  - Approaches are still evolving
  - Players are in and out everyday

IPTV *Must* Deliver Entertainment-Caliber Video

Tolerance is one visible artifact per movie

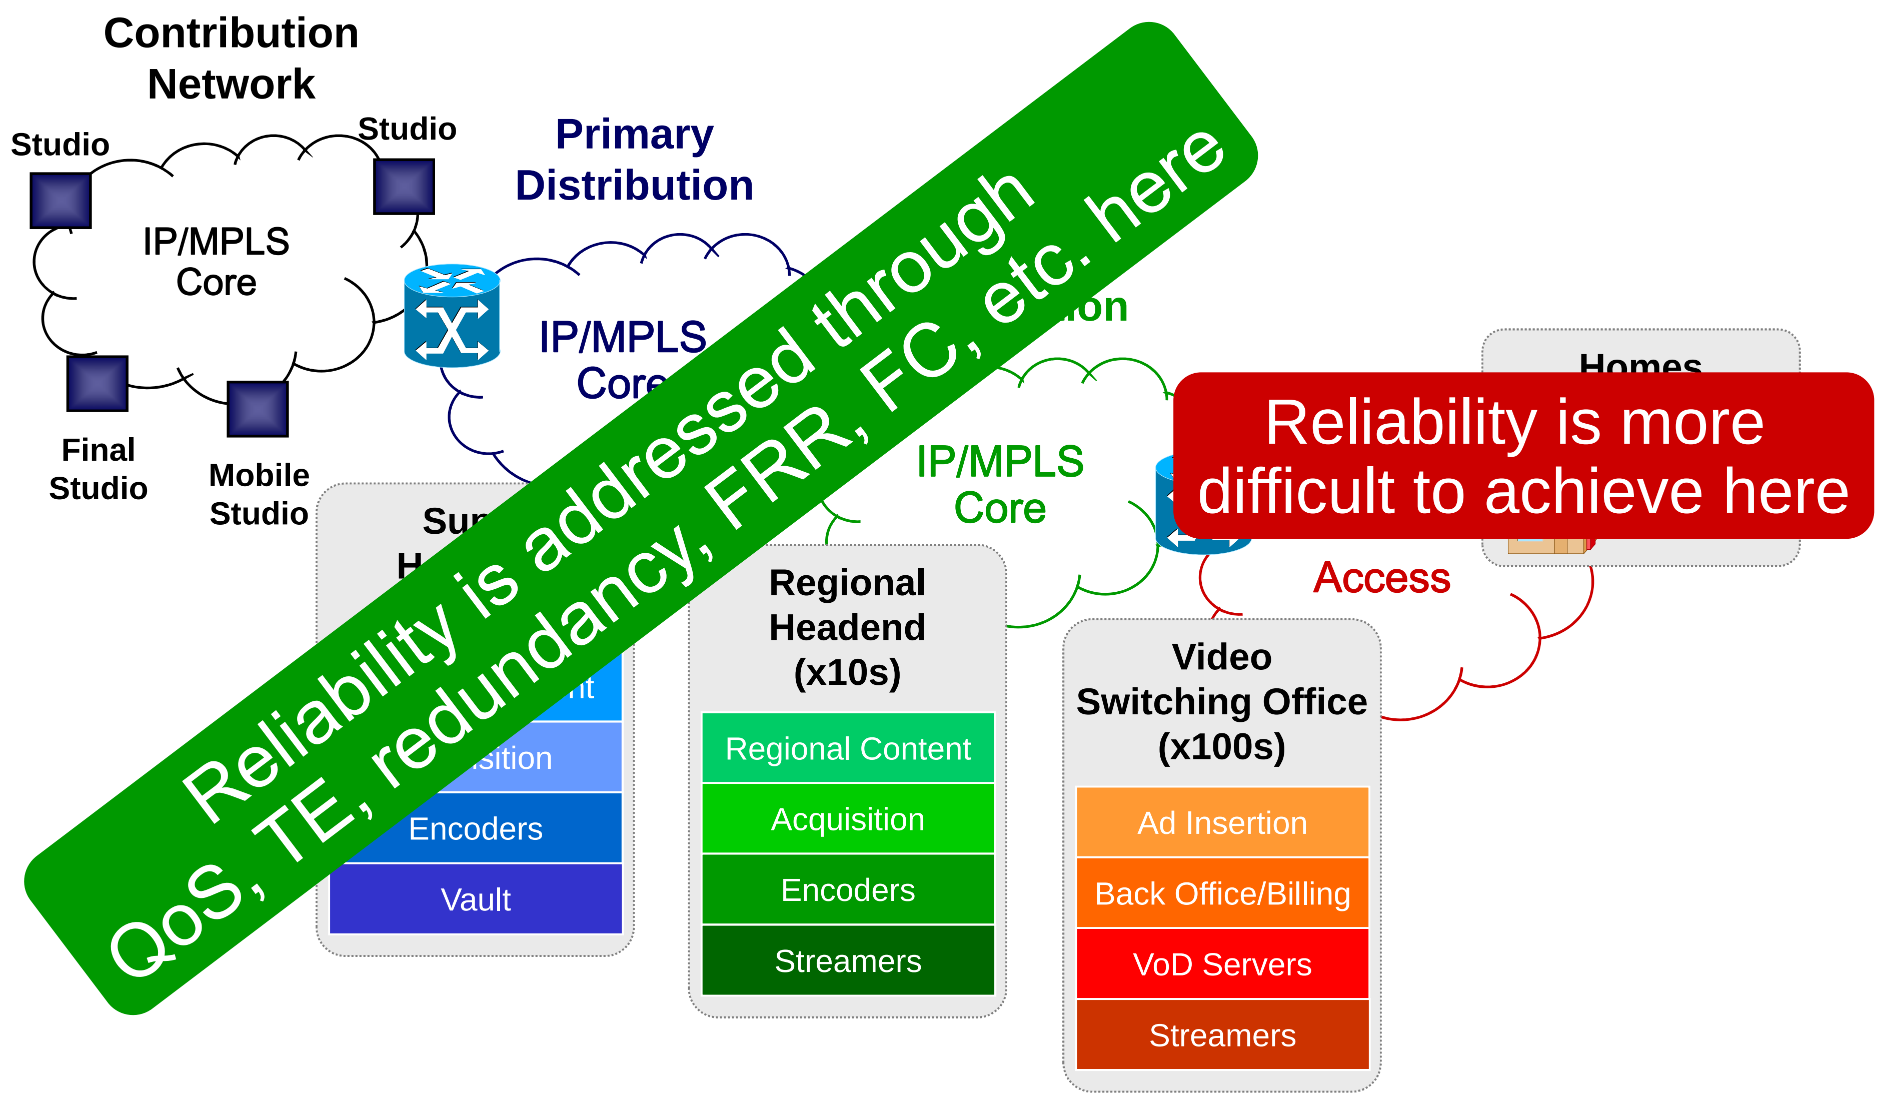

Taxonomy of Video Service Providers

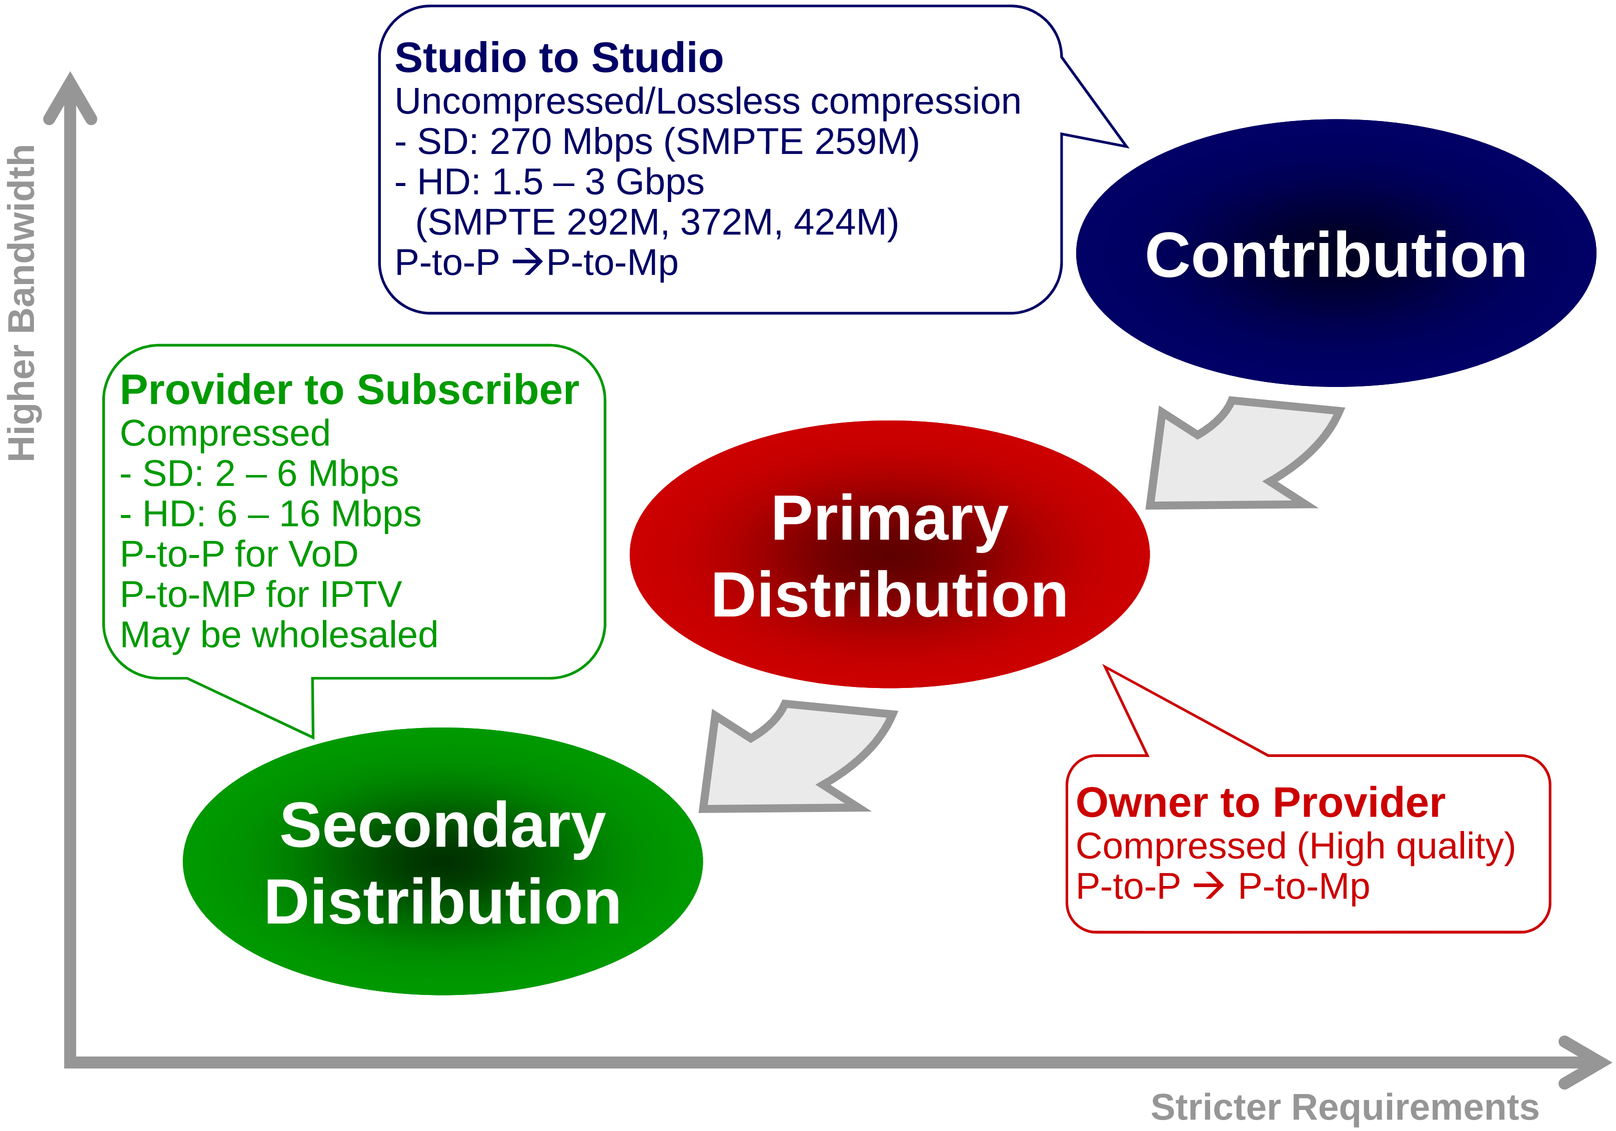

Digital Video Bandwidths

Uncompressed Digital Video	
SDTV (480i CCIR 601 over SD-SDI SMPTE 259M)	165.9 – 270 Mbps
EDTV (480p or 576p via SMPTE 344M)	540 Mbps
HDTV (1080i or 720p over HD-SDI SMPTE 292M)	1.485 Gbps
HDTV (1080p over Dual Link HD-SDI SMPTE 372M)	2.970 Gbps
MPEG-2 Compressed Video	
SDTV Broadcast (3.75 Mbps for cable VOD)	3 – 6 Mbps
HDTV Broadcast (19.3 Mbps for ATSC DTV)	12 – 20 Mbps
SDTV Production (Contribution – 4:2:2 I-frame only)	18 – 50 Mbps
HDTV Production (Contribution – 4:4:4 I-frame 10-bit)	140 – 500 Mbps
MPEG-4 AVC / H.264 Compressed Video	
SDTV Broadcast (~50% less than MPEG-2)	1.5 – 3 Mbps
HDTV Broadcast (1080i about 4x SDTV)	6 – 9 Mbps

Video SLA Requirements

- **Throughput**

    Addressed through capacity planning and QoS (i.e., Diffserv)

- **Delay/Jitter**

    Controlled with QoS

    Absorbed by de-jittering buffer at STB

    → We desire to minimize jitter buffer to improve responsivity

    → Jitter originating in the core is rather insignificant

- **Loss**

    Controlling loss is the main challenge

- **Service Availability**

    Proportion of time for which the specified throughput is available within the bounds of the defined delay and loss

    → A compound of the other networks and network availability

Video SLA Requirements

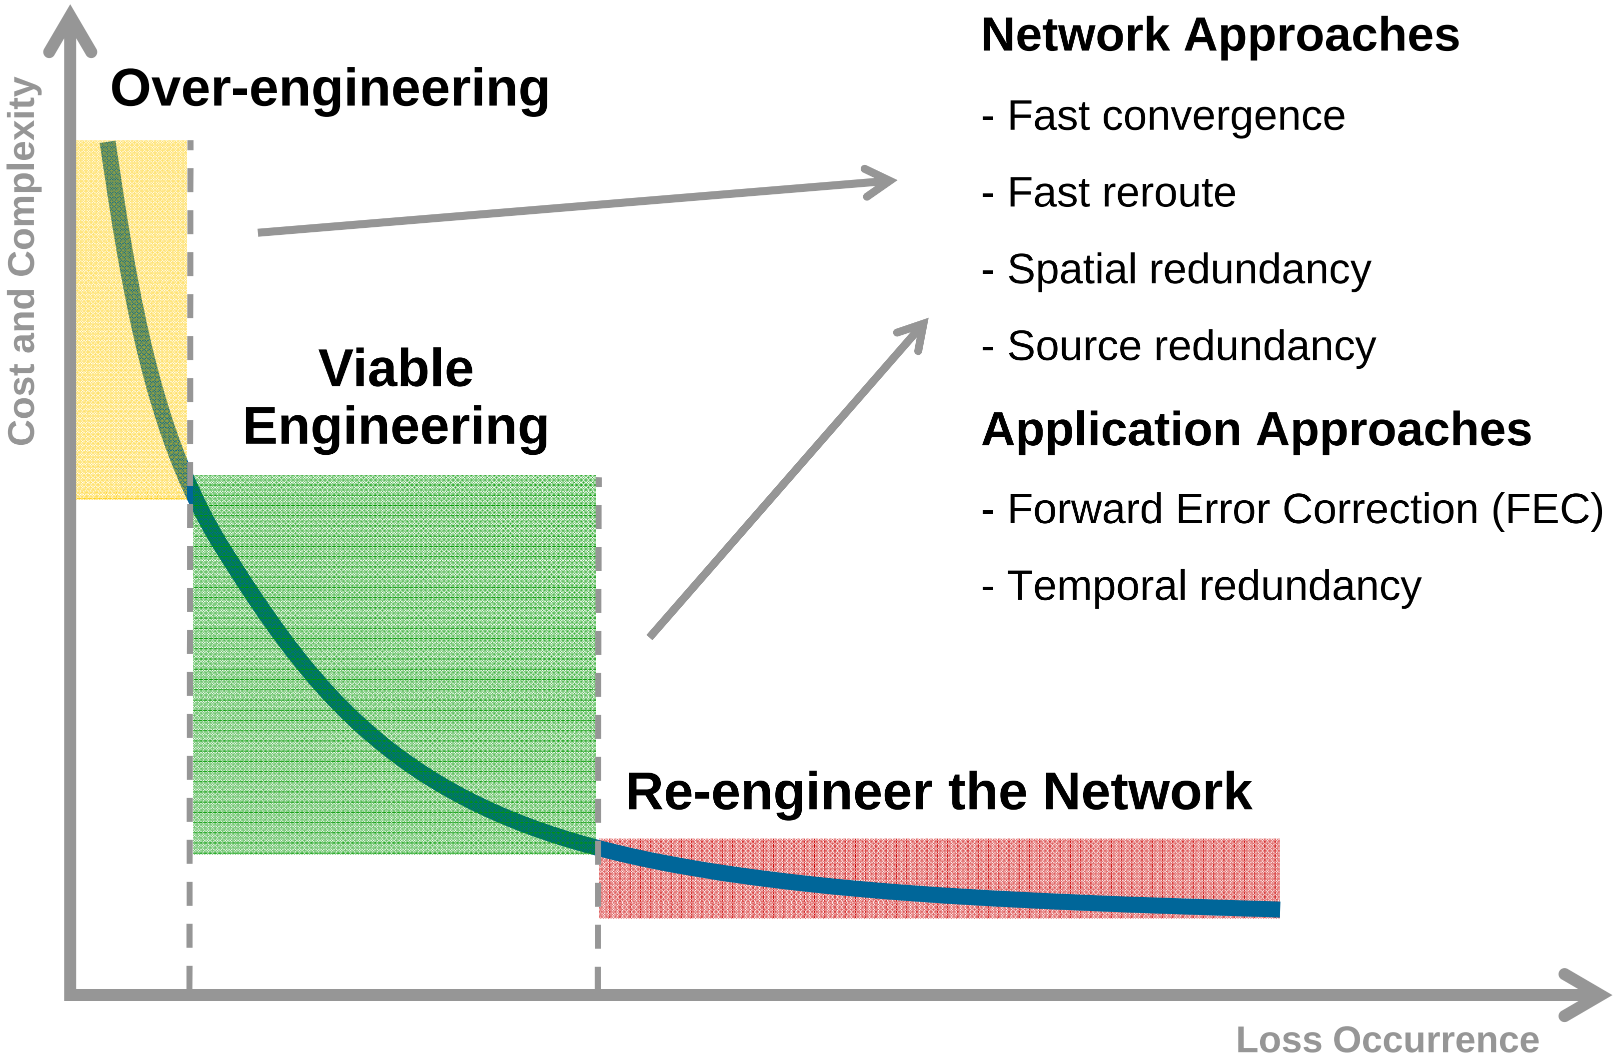

Network Approaches

- Fast convergence
- Fast reroute
- Spatial redundancy
- Source redundancy

Application Approaches

- Forward Error Correction (FEC)
- Temporal redundancy

Four Primary Causes for Packet Loss

- **Excess Delay**

  Renders media packets essentially lost beyond an acceptable bound

  Can be prevented with appropriate QoS (i.e., Diffserv)

- **Congestion**

  Considered as a catastrophic case, i.e., fundamental failure of service

  Must be prevented with appropriate QoS and admission control

- **PHY-Layer Errors (in the Core)**

  Apply to core and access – Occurrence in core is far less

  Assumed insignificant compared to losses due to network failures

- **Network Reconvergence Events**

  Occur at different scales based on topology, components and traffic

  Can be eliminated with high availability (HA) techniques

  → Impact of outage can be reduced with smart engineering

What are the Core Impairment Contributors?

	Impairment Rate
Trunk failures	**.0010 /2h**
Hardware failures	**.0003 /2h**
Software failures	**.0012 /2h**
Non-stop forwarding (NSF) and Stateful switch-over (SSO) help here	
Software upgrades (Maintenance)	**.0037 /2h**
Modular code (IOS-XR) helps here	
Total	**.0062 /2h** **(~One every two weeks)**

Note that average mean time between errors on a DSL line is in the order of minutes if no protection is applied

Back of envelope calculations across several SPs show mean time between core failures affecting video is > 100 hours

Based on assumptions, data from industry standards and customers

Il Buono, il Brutto, il Cattivo

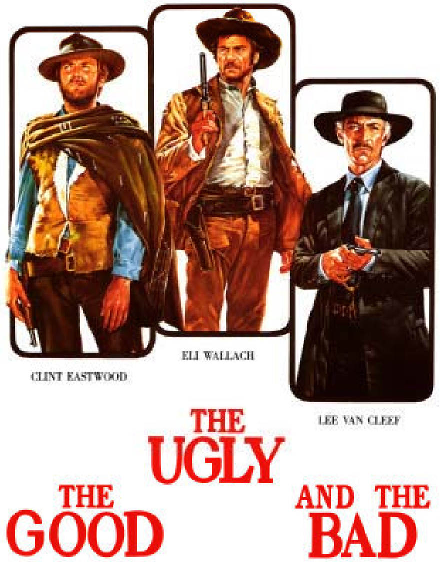

No Loss – Perfect Quality

0.5% Packet Loss

5% Packet Loss

MPEG Frame Impact from Packet Loss

GoP Size: 500 ms (I:P:B = 7:3:1)

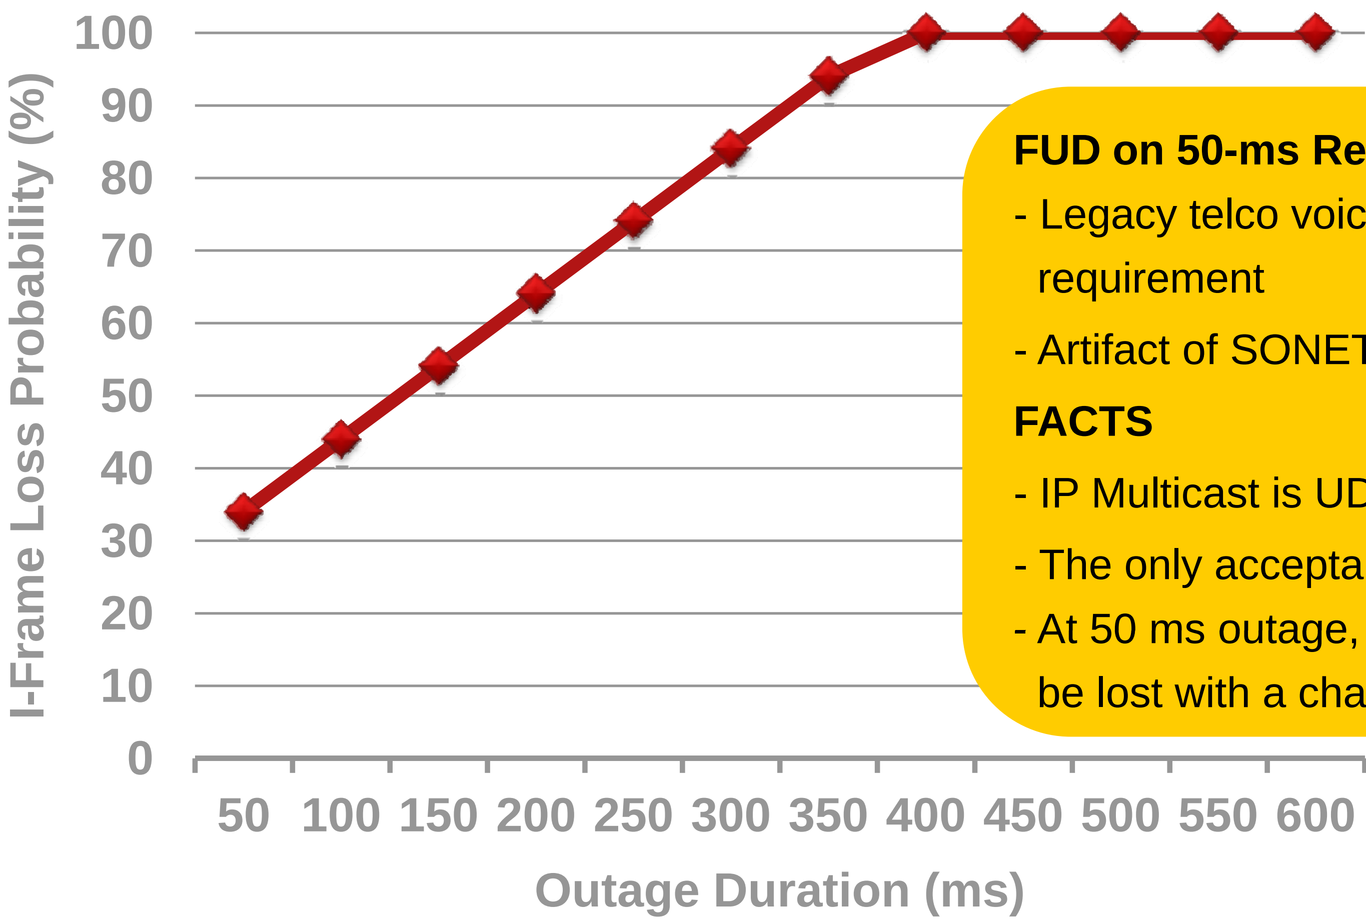

FUD on 50-ms Requirement

- Legacy telco voice (not video) requirement
- Artifact of SONET/SDH

FACTS

- IP Multicast is UDP based
- The only acceptable loss is 0%
- At 50 ms outage, an I-frame will be lost with a chance of 34%

Fast Convergence or Fast Reroute

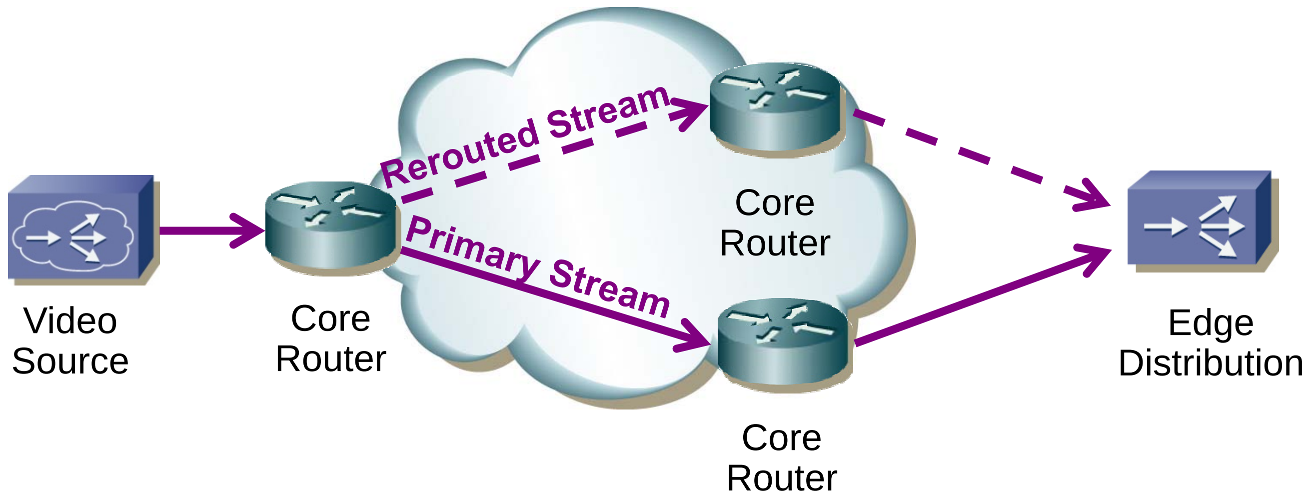

- Network reconverges / reroutes on core network failure (link or node)
- Fast convergence or Fast reroute
  - ✓ Lowest bandwidth requirements in working and failure case
  - ✓ Lowest solution cost and complexity
  - ! Requires fast converging network to minimize visible impact of loss
  - × Is NOT hitless – Loss of connectivity before connectivity is restored

Forward Error Correction (FEC)

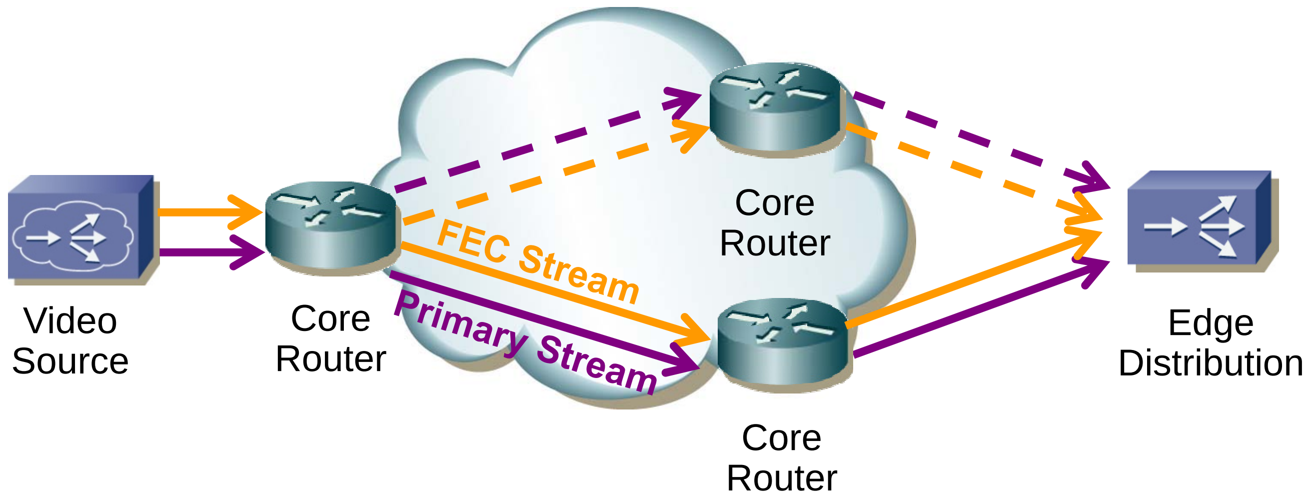

- FEC adds redundancy to the source data to allow the receiver to detect and repair errors (within some bound)
- FEC
  - ✓ Is hitless from loss due to core network failures if loss can be constrained
  - ✓ Does not require path diversity – Works for all topologies
  - ! Requires fast converging network to minimize FEC overhead
  - × Incurs delay – Longer outages require larger overhead or larger block sizes

Temporal Diversity

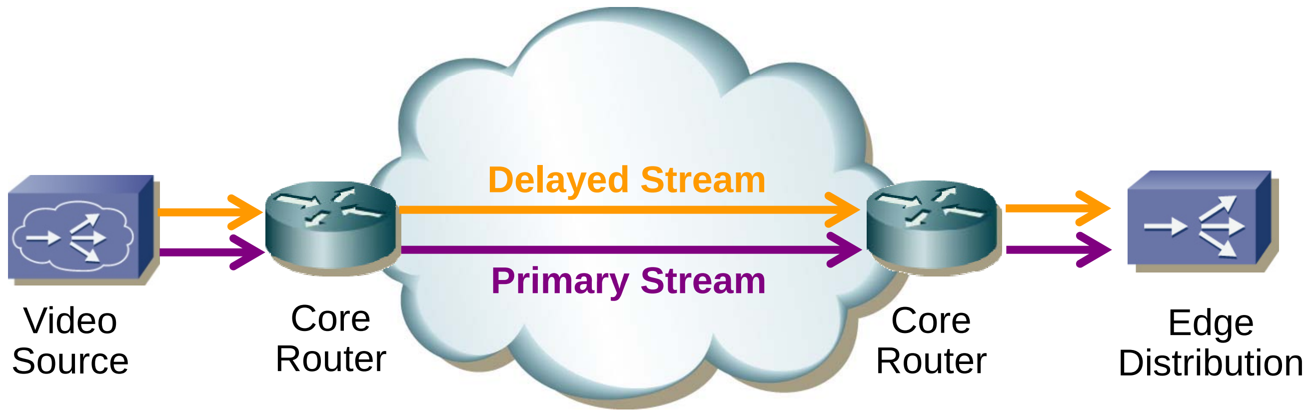

- Let Q (ms) denote the max outage duration that is intended to be repaired
- Packets are transmitted twice, each separated by Q-ms delay
- Temporal diversity
  - ✓ Is hitless from loss due to core network failures if loss can be constrained
  - ✓ Does not require path diversity – Works for all topologies
  - ! Requires fast converging network to minimize Q
  - × Introduces 100% overhead
  - × Introduces Q-ms delay

Spatial (Path) Diversity – Live/Live

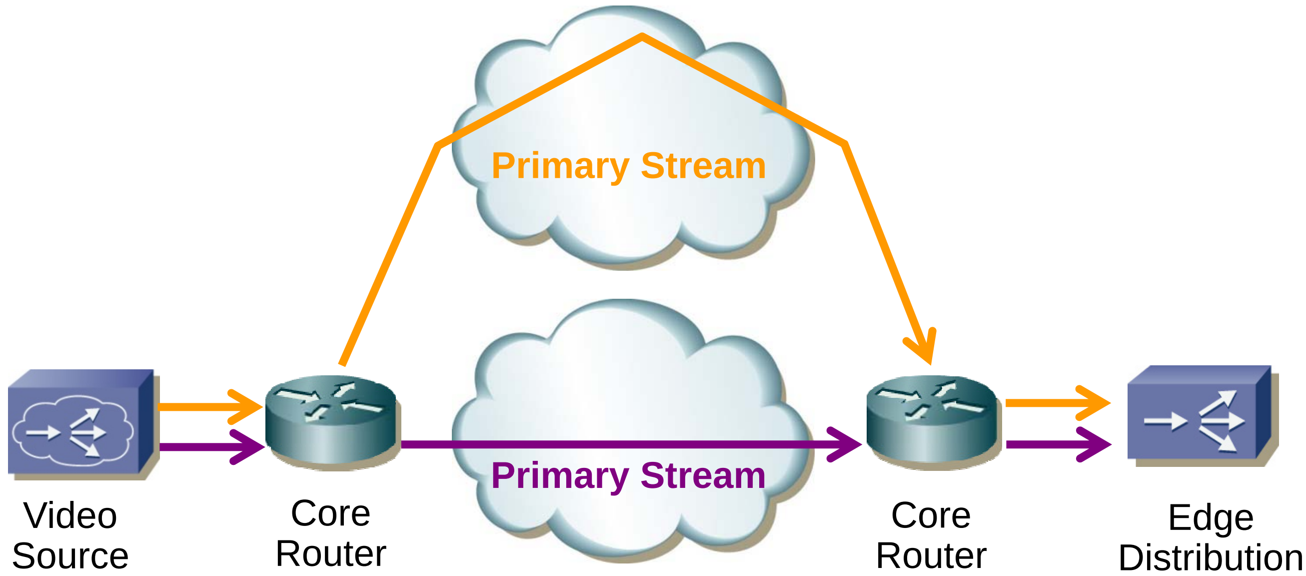

- Two streams are sent over diverse paths in the core
- Spatial (Path) diversity
  - ✓ Introduces no delay if the paths have equal propagation delays
  - × Requires network-level techniques to ensure spatial diversity
  - × Incurs 100% overhead
    - – May not be an issue where redundant capacity is normally provisioned
    - – E.g., dual-plane core networks

Source (Site) Diversity

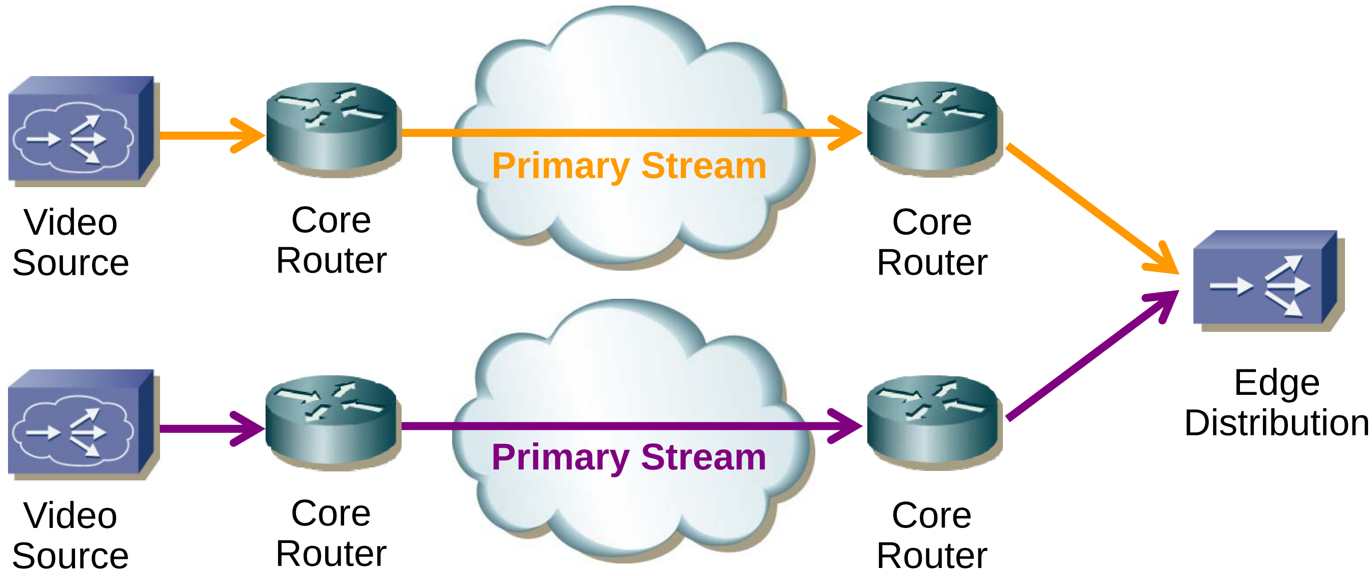

- Source (Site) diversity
  - ✓ Introduces no delay if the paths have equal propagation delays
  - ✓ May not require network-level techniques to ensure spatial diversity
    - – Topology dependent
  - × Incurs 100% overhead
    - – May not be an issue where redundant capacity is normally provisioned

This provides protection against single point-of-failure at the source, but this is **NOT a hitless** recovery as the sources are not in sync

Spatial Diversity through Multiple Interfaces

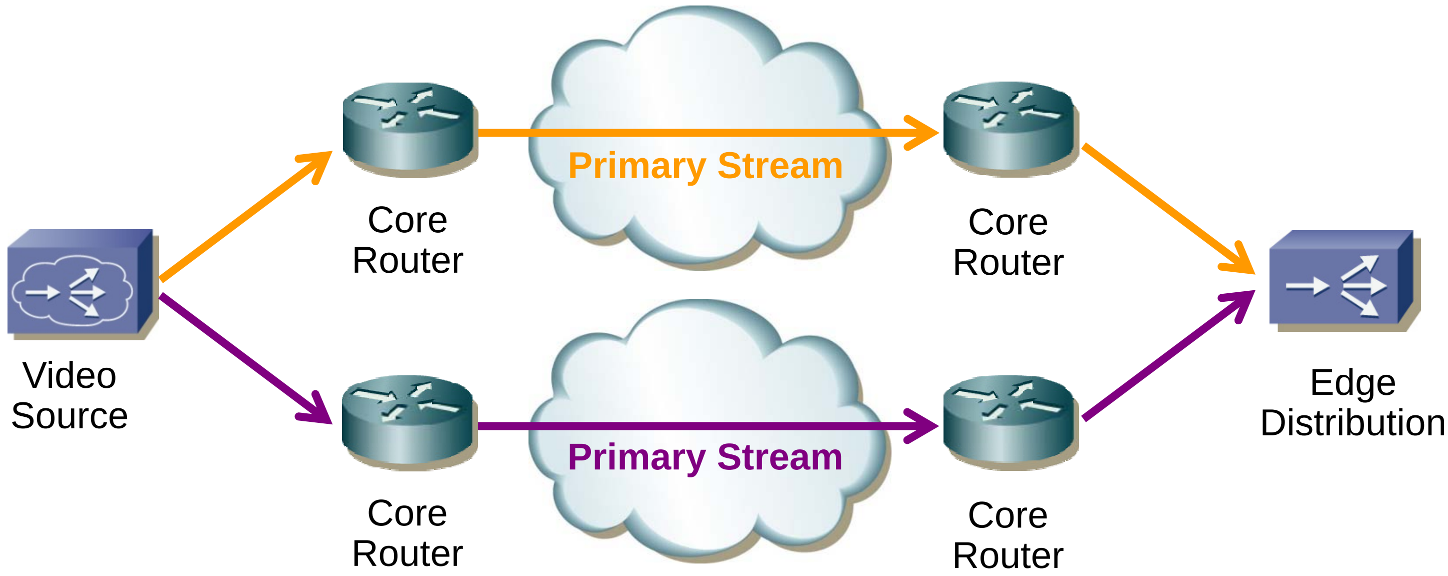

- Assume the headend is connected to two disjoint networks
- This approach
  - × Incurs 100% overhead
    - – May not be an issue where redundant capacity is normally provisioned
  - ✓ May not require network-level techniques to ensure spatial diversity
    - – Topology dependent
  - ✓ Offers a **hitless** recovery as the both primary streams are in sync

Towards Lossless IPTV Transport
Deployment Scenarios

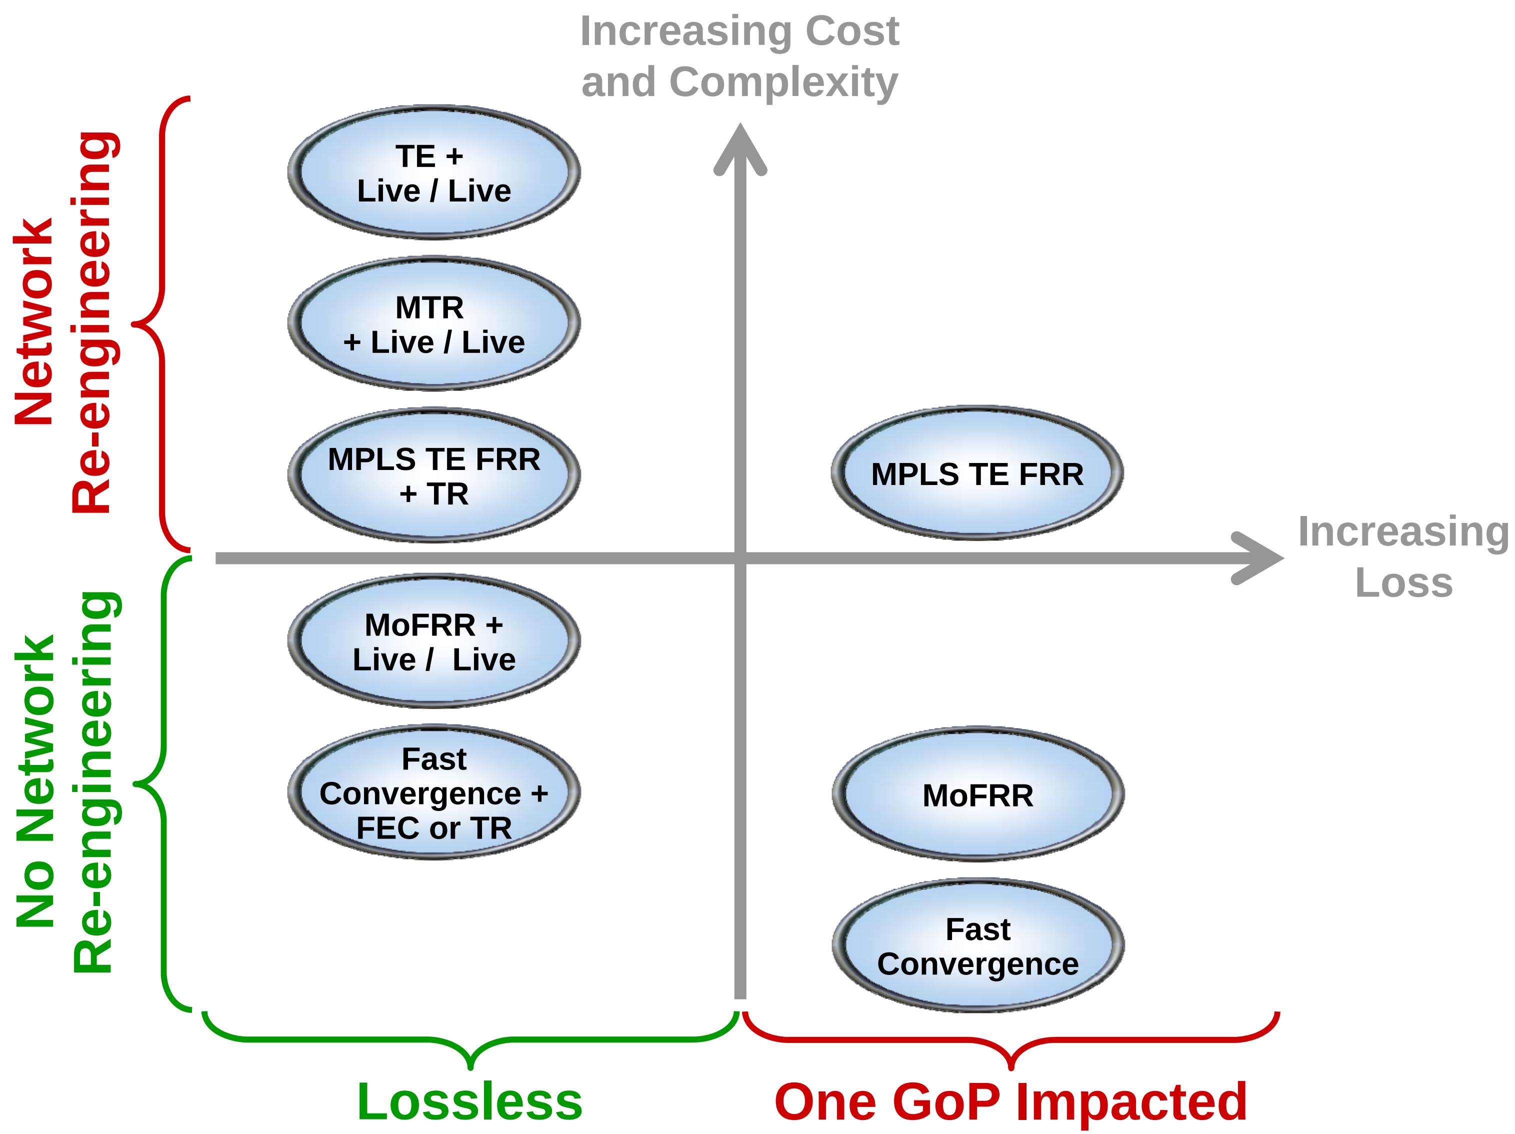

VQE – A Unified QoE Solution

- **IPTV viewers have two criteria to judge their service**
  - **Artifact-free audiovisual quality**
    - Packets dropped in access and home networks must be recovered quickly
    - Packet loss may or may not be correlated in spatial and/or temporal domain
    - Loss-repair methods must be multicast friendly
  - **Short and consistent zapping times**
    - Compression and encryption used in digital TV increase the zapping times
    - Multicasting in IPTV increases the zapping times
    - Zapping demand varies the zapping times
- **Service providers need a scalable unified solution that**
  - Is standards-based and interoperable with their infrastructure
  - Enables versatility, quick deployment and visibility into the network
  - Extends the service coverage area, and keeps CAPEX and OPEX low
- **Our goals are to offer**
  - Glitch-free audiovisual quality, short and consistent zappings even in low-bandwidth networks
  - Monitoring tools that isolate and pinpoint the problematic locations

VQE does for video what Dolby did for stereo

Real-Time Transport Protocol (RTP)

- **Basics**
  - First specified by IETF in 1996, later updated in 2003 (RFC 3550)
  - Runs over any transport-layer protocol – UDP is much more widely used
  - Runs over both unicast and multicast
  - No built-in reliability
- **Main Services**
  - Payload type identification
  - Sequence numbering
  - Timestamping
- **Extensions**
  - Basic RTP functionality uses a 12-byte header
  - RFC 5285 defines an RTP header extension mechanism
- **Control Plane – RTCP**
  - Provides minimal control and identification functionality
  - Enables a scalable monitoring functionality (Sender, receiver and extended reports)
- **RTP Transport**
  - Terrestrial, satellite and emerging IPTV networks dominantly use MPEG2-TS encapsulation
  - RFC 2250 defines a way to carry TS packets within RTP packets

A Simplified Model

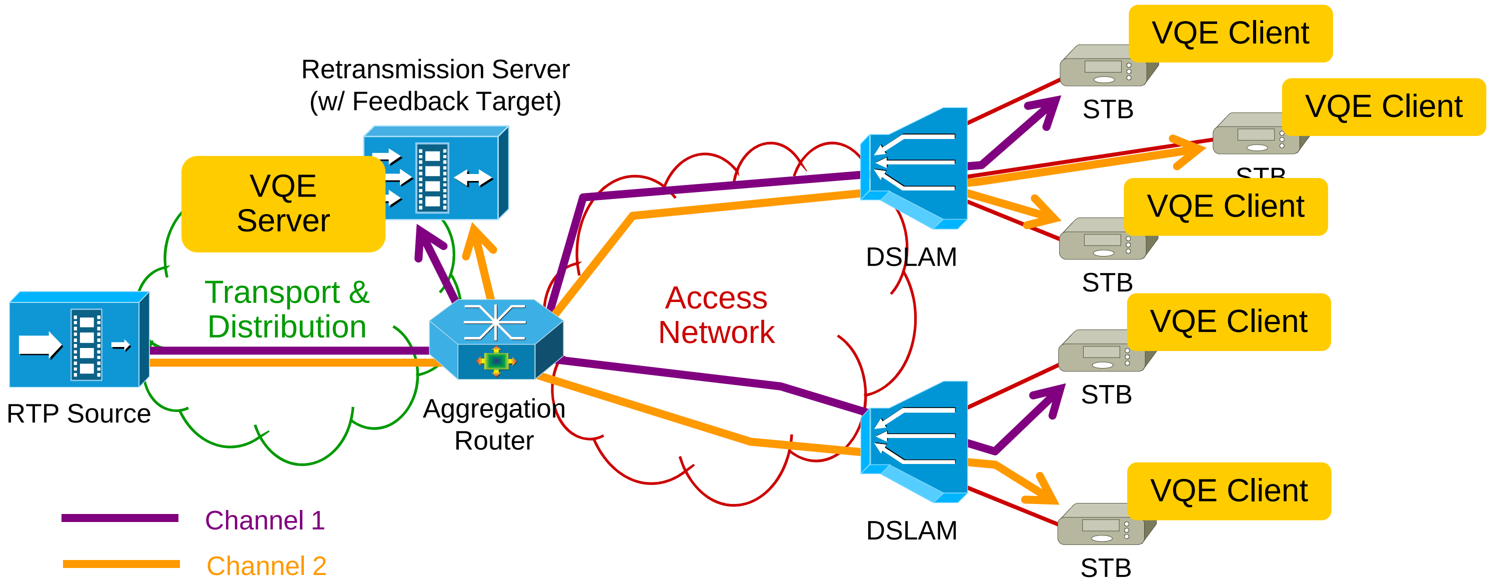

- Each TV channel is served in a unique (SSM) multicast session
  - IP STBs join the respective multicast sessions for the desired TV channel
  - Retransmission servers join all the multicast sessions
- (Unicast) Feedback from IP STBs are collected by the feedback target
  - NACK messages reporting missing packets
  - Rapid channel change requests
  - RTCP receiver and extended report reports reporting reception quality

Packet Loss Rate Tolerance Limits

Each random or bursty loss counts for one artifact

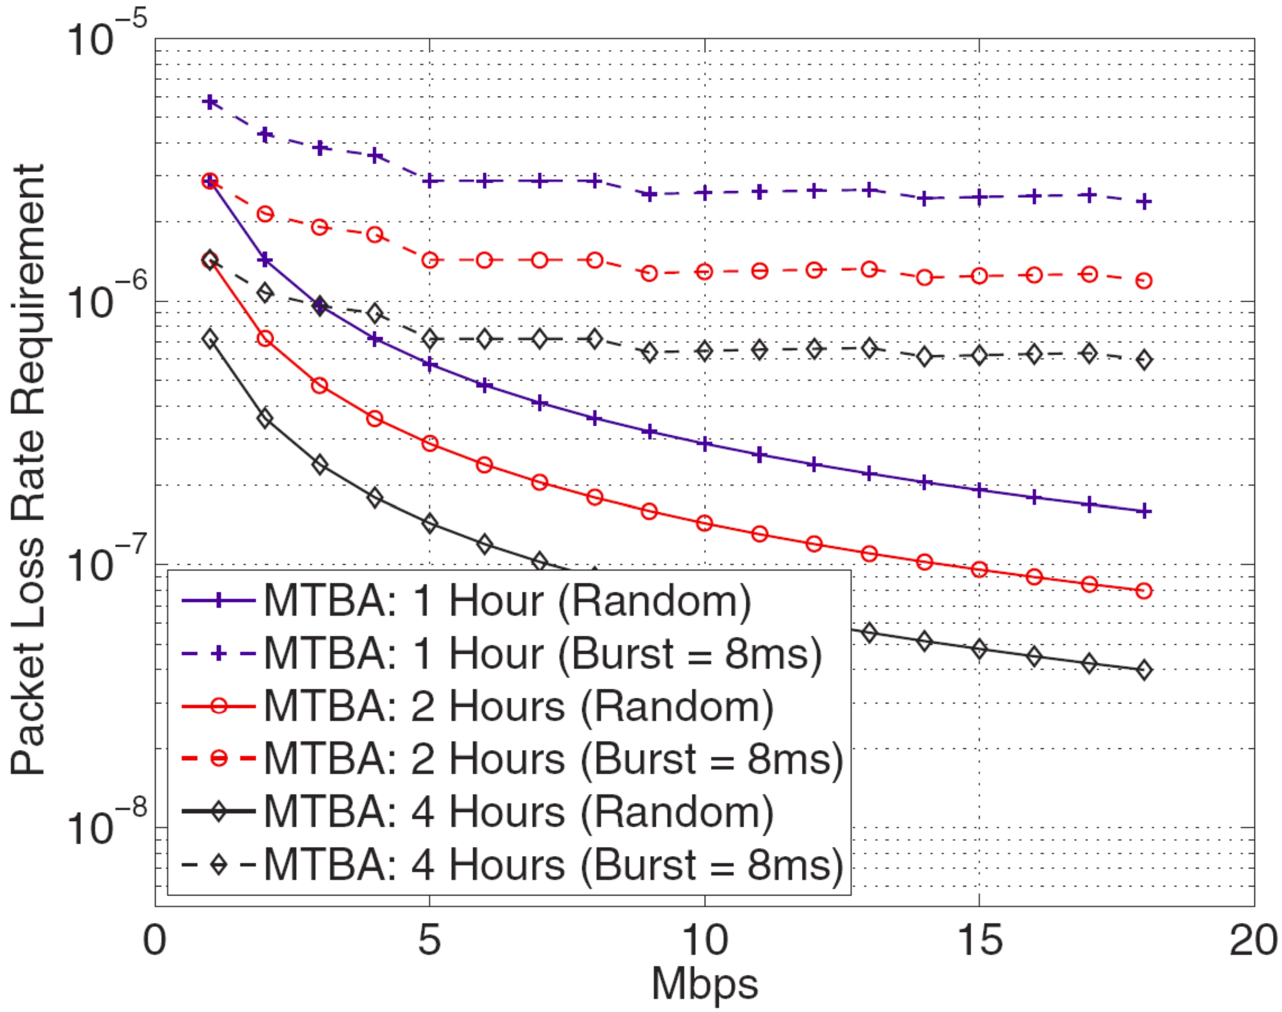

Impairments in xDSL Networks

- Twisted pair is subject to
  - Signal attenuation → Use shorter loops
  - Cross talk → Use Trellis Coding and RS-based FEC
  - Impulse noise → Use RS-based FEC with interleaving
- Three types of DSL impulse noise
  - REIN → Short burst of noises (< 1 ms)
  - PEIN → Individual impulse noise (> 1 ms, < 10 ms)
  - SHINE → Individual impulse noise (> 10 ms)
- We observe different noise characteristics
  - Among different SP networks
  - Among different loops in the same SP network

ADSL Transmitter Reference Model

ITU-T Recommendation G.992.1

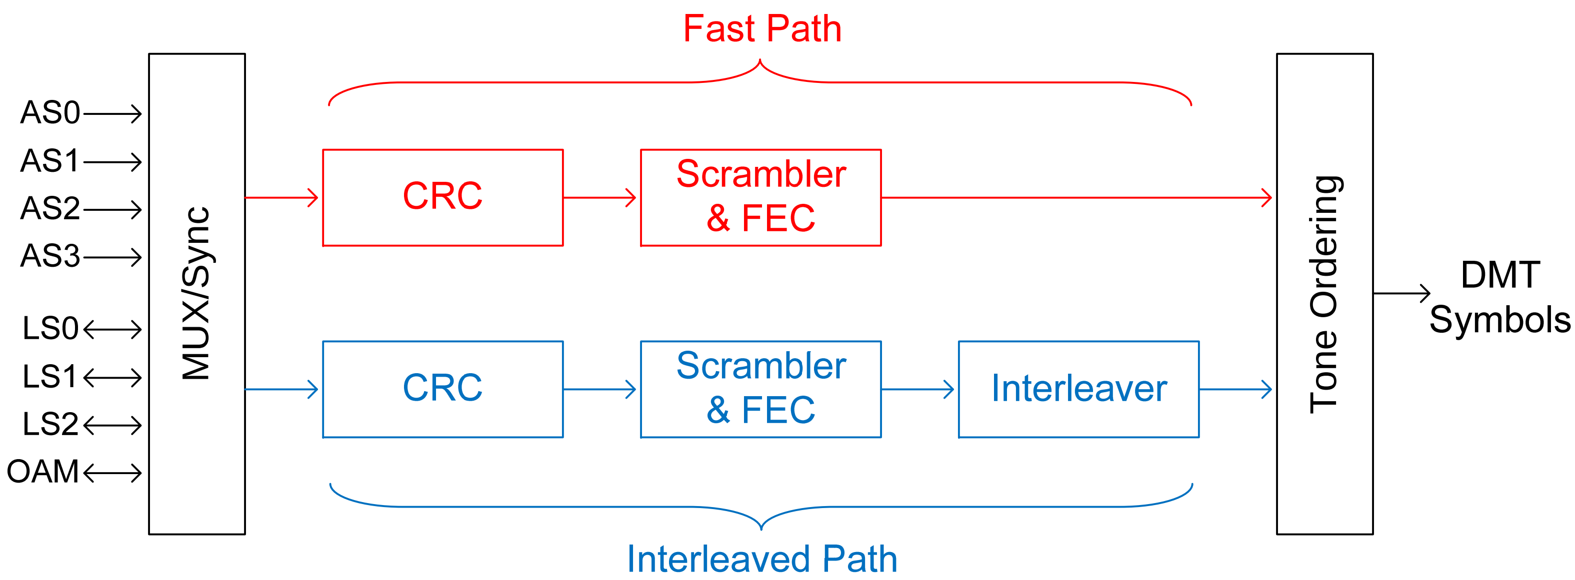

ADSL and ADSL2+ Configurations

	ADSL	ADSL2+
Data bytes per RS codeword, K	239 bytes	69 bytes
Parity bytes per RS codeword, R	16 bytes	10 bytes
Correctable byte errors per RS codeword, T = R / 2	8 bytes	5 bytes
Total bytes per RS codeword, N = K + R	255 bytes	79 bytes
# of RS codewords per DMT symbol, 1/S	1	11
DMT duration, t	250 us	250 us
Line data rate, LDR = N / S / t	8.0 Mbps	27.4 Mbps
Net data rate, NDR = LDR x K / N	7.5 Mbps	24 Mbps
Interleaver depth, D	32	352
Size of required memory, B = (N–1) x (D–1)	7874 bytes	27378 bytes
Interleaving delay, ID = B / LDR	7.87 ms	7.97 ms
Block size (Protection period), PP = N x D / LDR	8.16 ms	8.10 ms
Correctable error burst length, BL = D x T	256 bytes	1760 bytes
Impulse noise protection, INP = floor(BL / (N / S))	1	2

Example: Interleaving of RS Codewords

Original RS Codewords

Interleaved RS Codewords

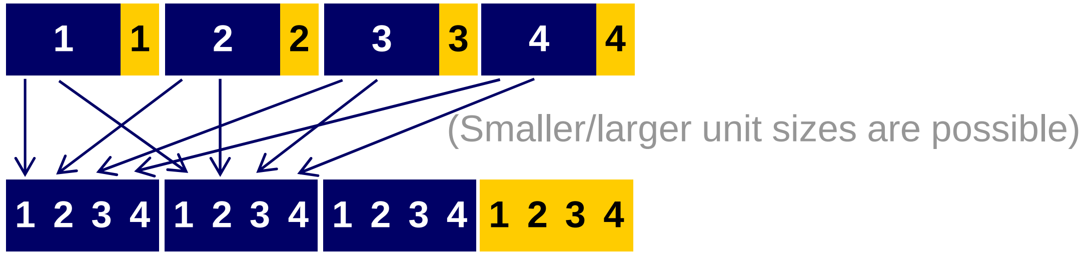

- Interleaving
  - ✓ Spreads a bursty error among multiple codewords
  - ✓ Allows the decoder to repair the error with fewer parity bytes
  - × Introduces delay
  - × Renders the whole block useless upon a decoding failure

Fast vs. Interleaved Path

- Assumptions

  One impulse noise arrives every 15 seconds

  2% of these impulses cause an error

- Conditional probability of (DMT error | There is an error)

  85% → One DMT in error

  12% → Two DMTs in error

  3% → Three or more DMTs in error

- Fast Path

  ADSL/ADSL2+: One (maybe two) IP packet loss in every 750 seconds

- Interleaved Path (Interleaving delay: 8ms)

  ADSL: Up to 7 IP packet losses (at the net rate) in every 5000 seconds

  ADSL2+: Up to 19 IP packet losses in every 25000 seconds

First-Line of Defense in Loss Repair

1-D/2-D Parity Forward Error Correction

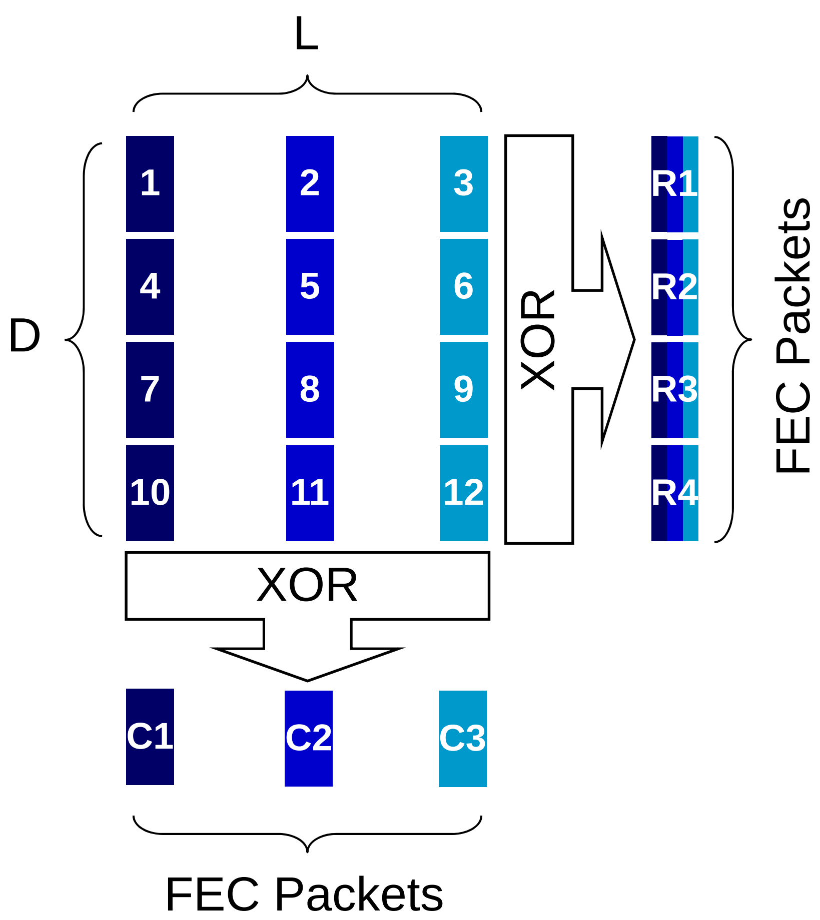

- Source block size: D x L
- 1-D Column FEC (for Bursty Losses)

  Each column produces a single packet

  Overhead = 1 / D

  L-packet duration should be larger than the (target) burst duration
- 1-D Row FEC (for Random Losses)

  Each row produces a single packet

  Overhead = 1 / L
- 2-D Column + Row FEC

  Overhead = (D+L)/(DxL)

First-Line of Defense in Loss Repair
1-D/2-D Parity Forward Error Correction

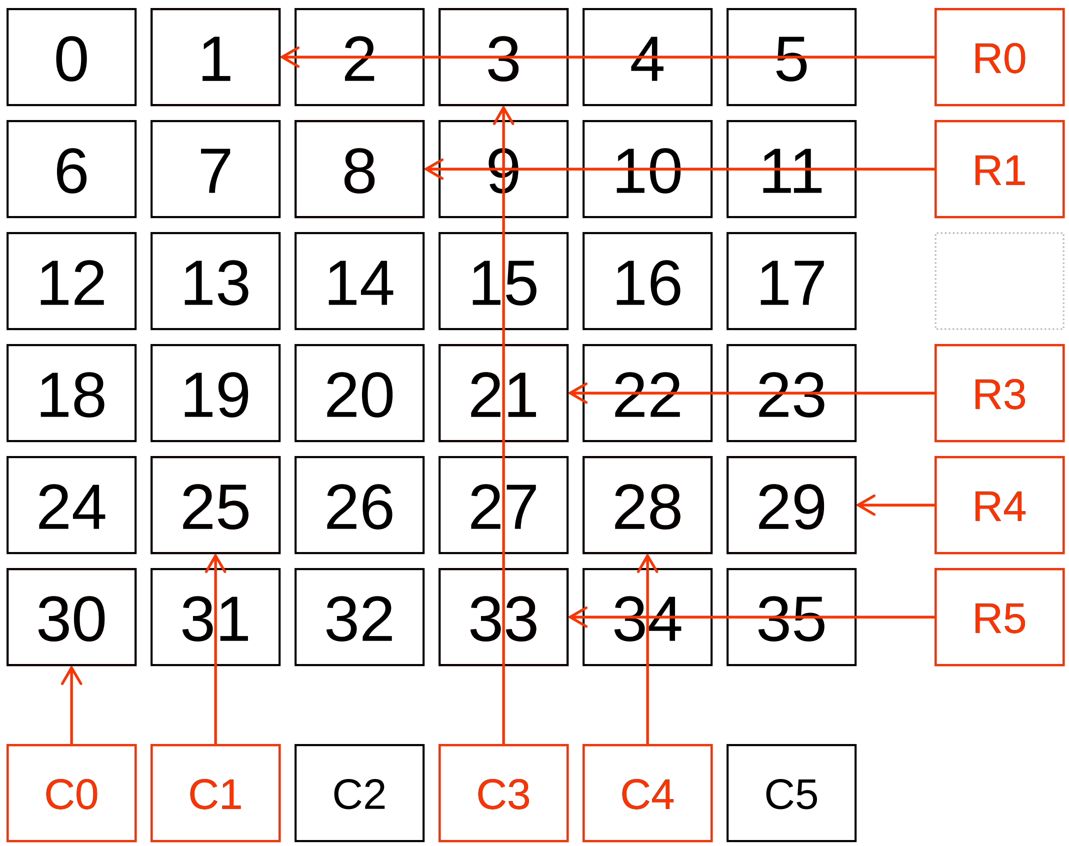

All nine missing data packets are successfully recovered

First-Line of Defense in Loss Repair
1-D/2-D Parity Forward Error Correction

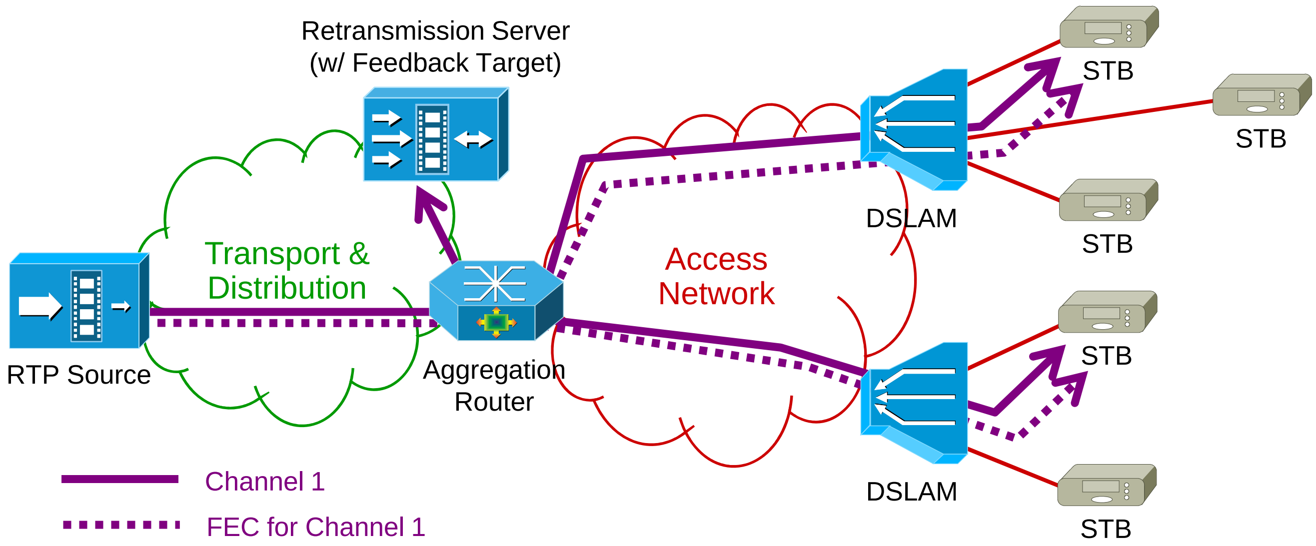

- Each TV channel may be associated with one or more FEC streams
  - FEC streams may have different repair capabilities
  - IP STBs may join the respective multicast sessions to receive FEC stream(s)
- General Remarks
  - ✓ FEC scales extremely well with upfront planning, easily repairs spatially correlated losses
  - × Longer outages require larger overhead or larger block sizes (more delay)
  - × FEC requires encoding/decoding operations

Second-Line of Defense in Loss Repair
RTP Retransmissions

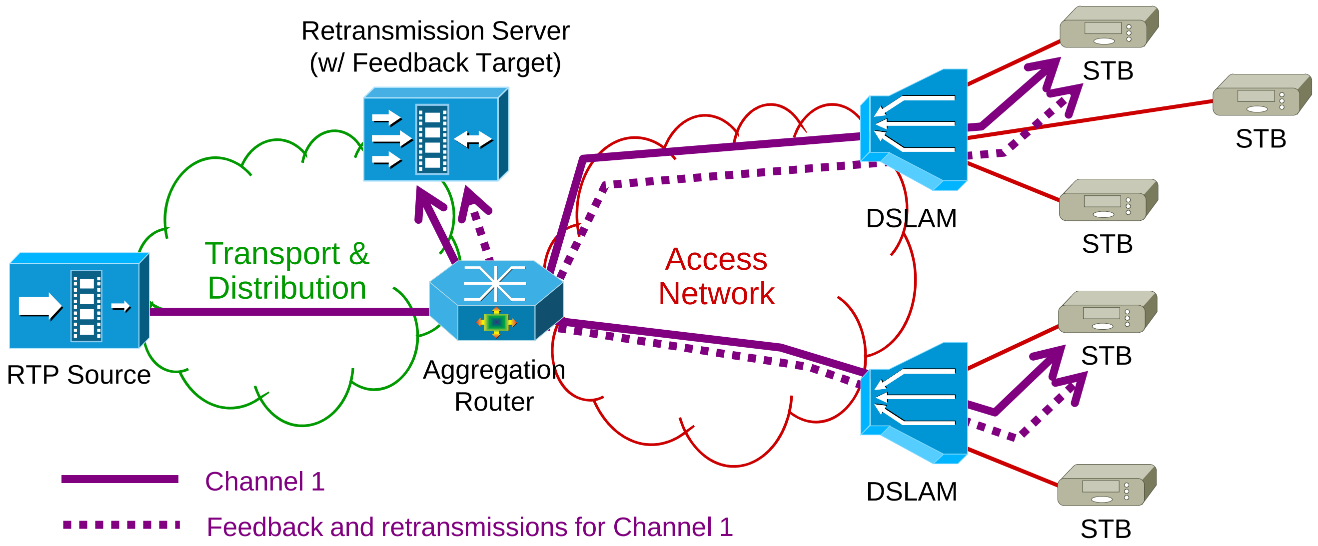

- There is a (logical) feedback target for each TV channel on the retransmission server
  - If optional FEC cannot repair missing packets, STB sends an RTCP NACK to report missing packets
  - Retransmission server pulls requested packets out of the cache and retransmits them
  - The retransmission is on a separate unicast RTP session
- General Remarks
  - ✓ Retransmission recovers only the lost packets, so no bandwidth is wasted
  - × Retransmission adds a delay of destination-to-source-to-destination
- Protocol suite comprises RFC 3550, 4585, 4588 and RTCP SSM

TV Viewers Love Zapping

Results are based on 227K+ users in NA

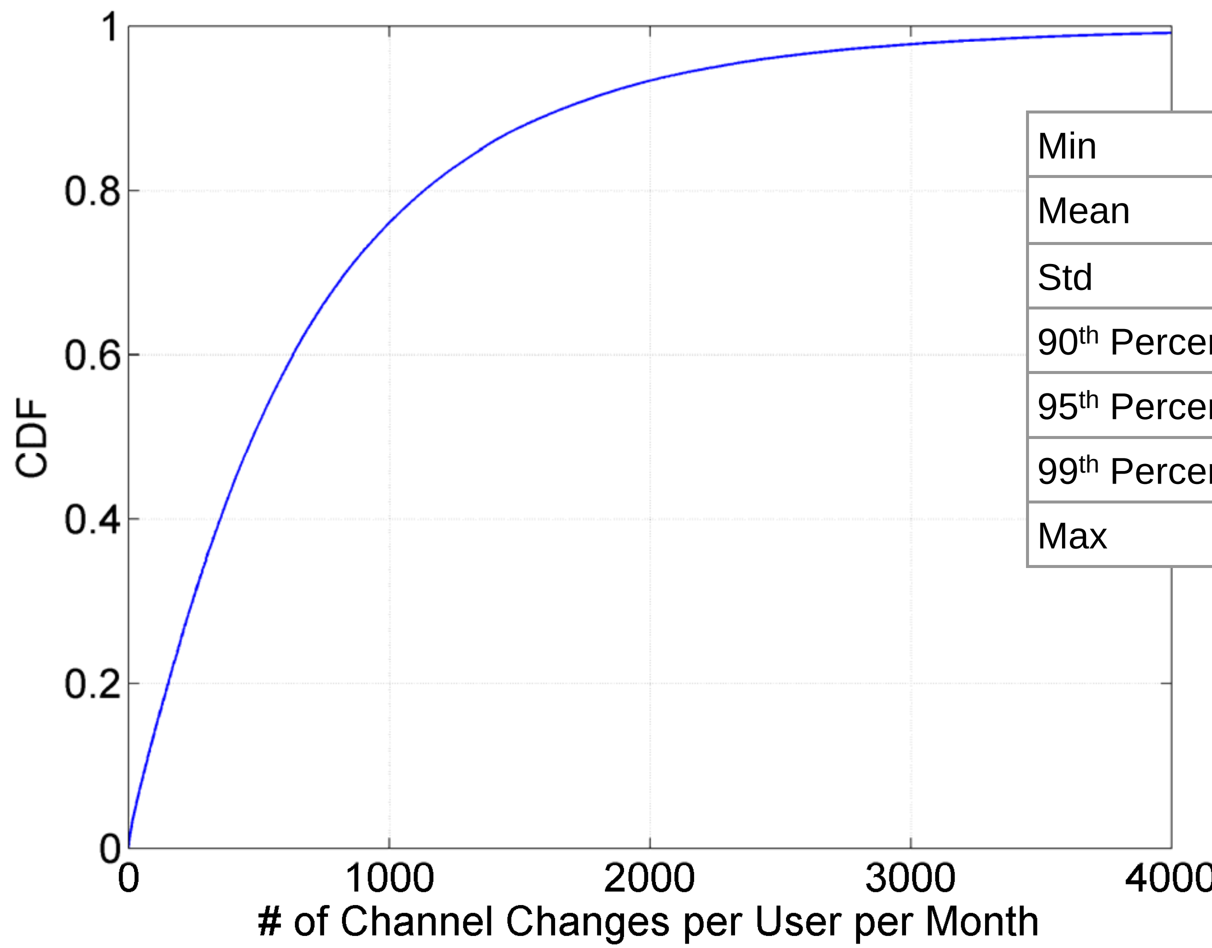

Min	1
Mean	726
Std	814
90^{th} Percentile	1672
95^{th} Percentile	2250
99^{th} Percentile	3798
Max	24186

Zappings are Correlated in Temporal Domain

On a Sunday between 8:00 – 9:00 PM

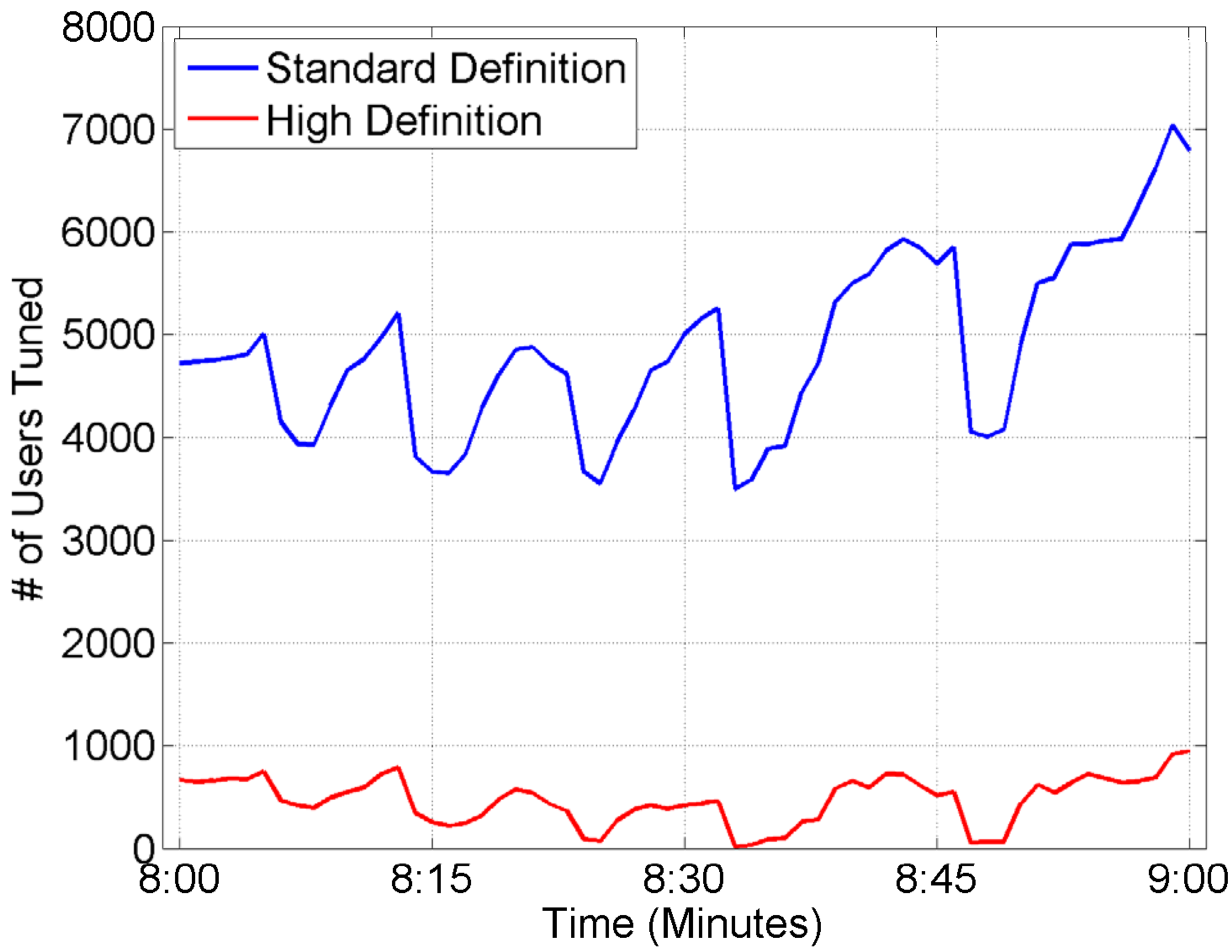

Delay Elements in Multicast MPEG2-TS Video

- **Multicast Switching Delay**
  - IGMP joins and leaves
  - Route establishment (Generally well-bounded)
- **Reference Information Latency**
  - PSI (PAT/CAT/PMT) acquisition delay
  - CAS (ECM) delay
  - RAP acquisition delay
- **Buffering Delays**
  - **Loss-repair, de-jittering, application buffering**
  - **MPEG decoder buffering**

Reference information latency and buffering delays are more critical in MPEG-based AV applications

Typical Zapping Times on DSL IPTV

	Unit Time	Total Time
STB sends IGMP Leave	**< 10 ms**	
STB sends IGMP Join	**< 10 ms**	
DSLAM gets IGMP Leave	**< 10 ms**	
DSLAM gets IGMP Join	**< 10 ms**	**~ 20 ms**
DSLAM switches streams	**30 ms**	**~ 50 ms**
Latency on DSL line	**~ 10 ms**	**~ 60 ms**
STB receives PAT/PMT	**~ 125 ms**	**~ 185 ms**
Buffering		
De-jittering buffer	**< 50 ms**	**~ 200 ms**
Wait for CA	**< 50 ms**	**~ 250 ms**
Wait for I-frame	**0 – 3 s**	**0.2 – 3.2 s**
MPEG decoding buffer	**1 – 2 s**	**1.2 – 5.2 s**
Decoding	**< 50 ms**	**1.2 – 5.2 s**

A Typical Multicast Join

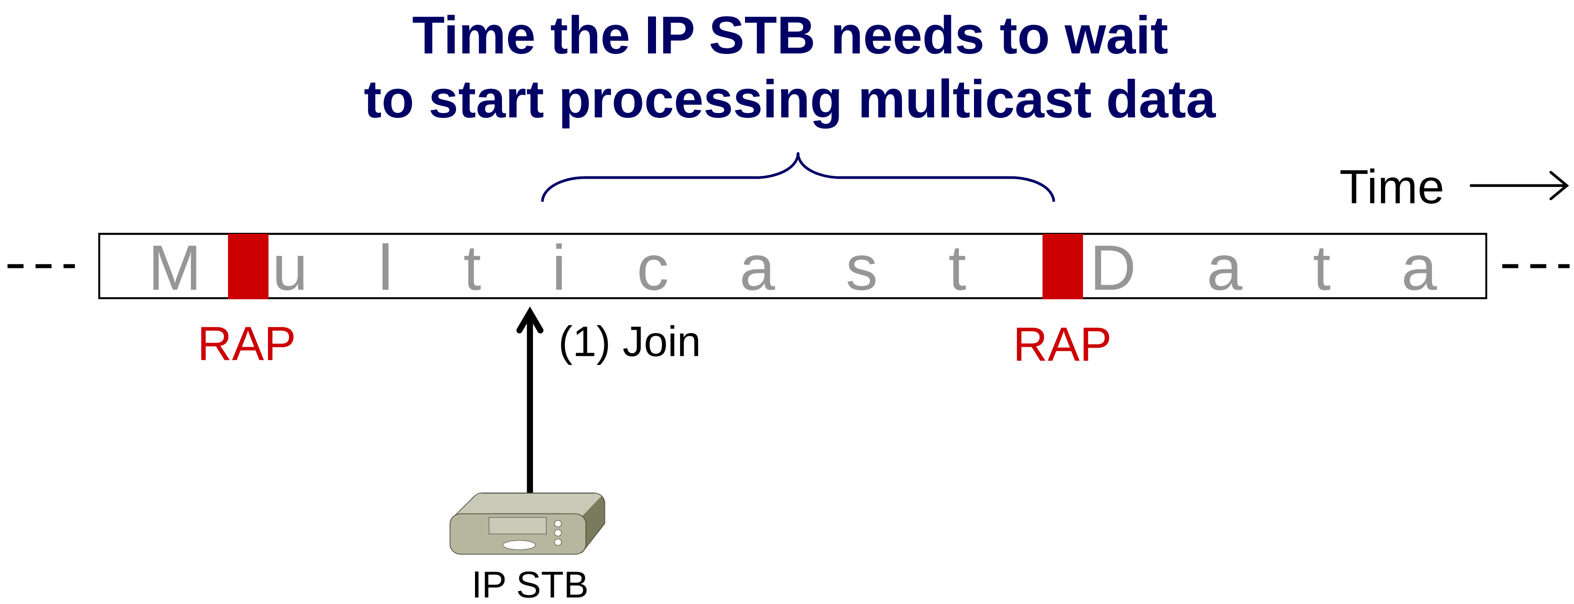

RAPs might be far away from each other
RAP data might be large in size and non-contiguous

Concurrent Multicast Join and Retransmission

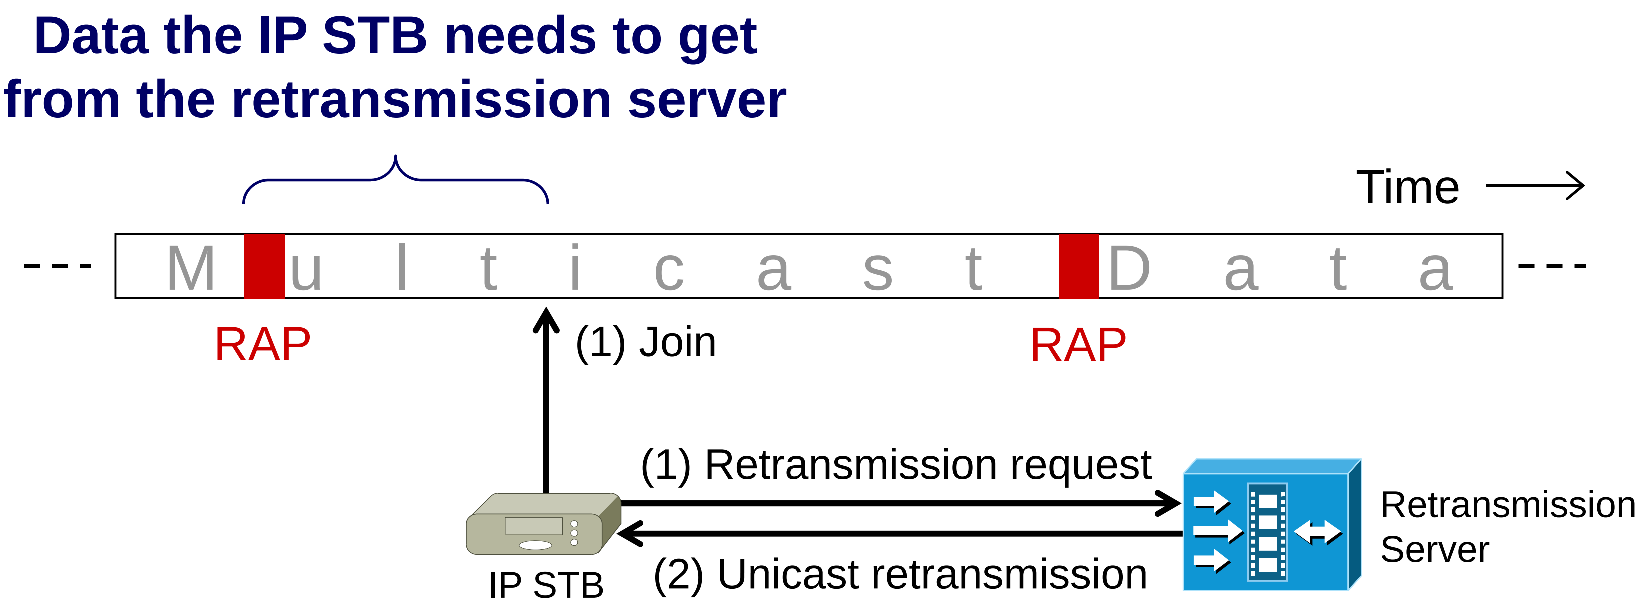

If the residual bandwidth remaining from the multicast flow is small, retransmission may not be able to provide acceleration

Retransmission Followed by Multicast Join

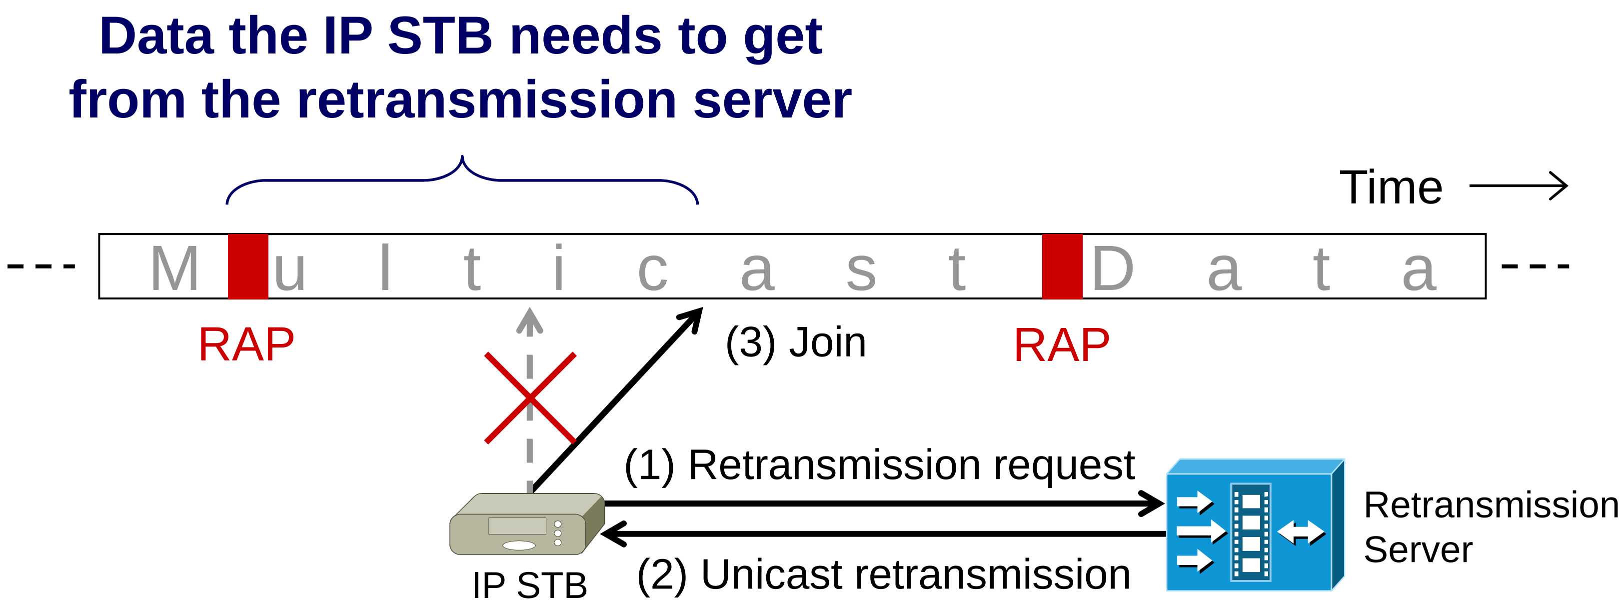

More data are retransmitted due to deferred multicast join
However, IP STB ultimately achieves a faster synchronization

Proposed Solution

Unicast-Based Rapid Acquisition of Multicast RTP Sessions

- IP STB says to the retransmission server:

“I have no synch with the stream. Send me a repair burst that will get me back on the track with the multicast session”

- Retransmission server may need to
  - Parse data from earlier in the stream than it is needed for retransmission
  - Burst faster than real time
  - Coordinate the time for multicast join and ending the burst
- This solution
  - Is applicable to any RTP-encapsulated multicast flow
  - Uses the existing toolkit for repairing packet losses in multicast sessions
    - RFC 3550 (RTP/RTCP)
    - RFC 4585 (RTP/AVPF)
    - RFC 4588 (RTP Retransmissions)
    - RTCP SSM (RTCP Extensions for SSM – *with the RFC Editor*)

Rapid Acquisition of Multicast Sessions

http://tools.ietf.org/html/draft-ietf-avt-rapid-acquisition-for-rtp

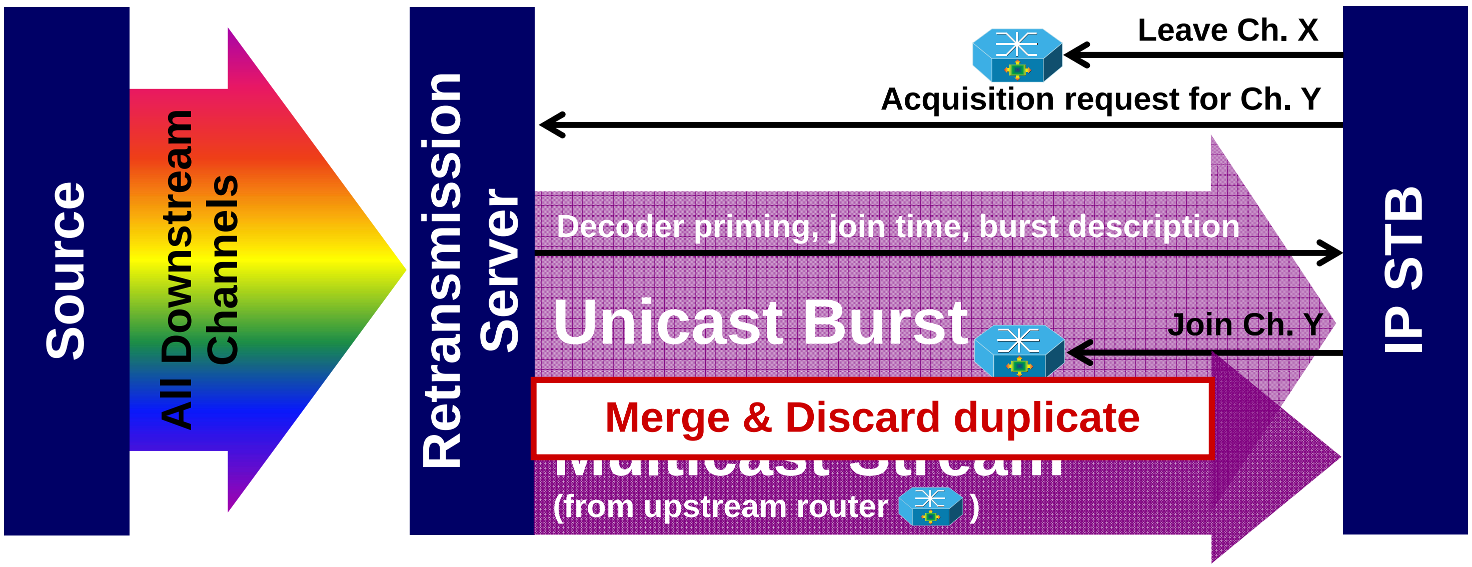

Retransmission server subscribes to all downstream multicast sessions

How to Prime the MPEG Decoder?

http://tools.ietf.org/html/draft-begen-avt-rtp-mpeg2ts-preamble

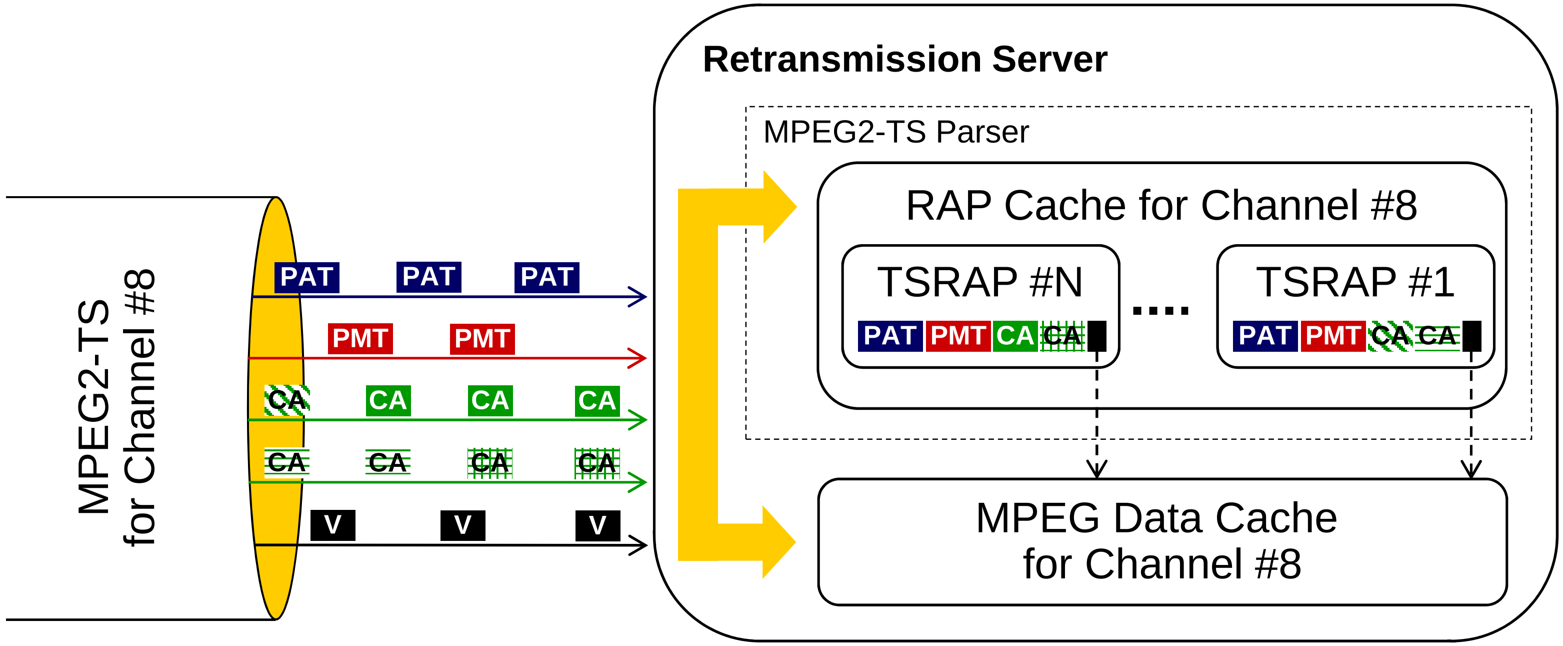

- Transport Stream Random Access Point (TSRAP) may include

  PAT: Program Association Table

  PMT: Program Map Table

  PCR: Program Clock Reference used to initialize the decoder and STB clocks

  SEQ: Sequence Header (MPEG2 video)

  SPS: Sequence Parameter Set (H.264 video)

  PPS: Picture Parameter Set (H.264 video)

  ECM: Entitlement Control Messages

Experimental Setup

- **Comparison**

  One IP STB with non-accelerated channel changes

  One IP STB with accelerated channel changes

- **Video Streams**

  High-detail, high-motion scenes of a movie

  AVC encoded at 2 Mbps and 30 fps

    One stream with 15 frames per GoP (Short-GoP)

    One stream with 60 frames per GoP (Long-GoP)

- **Transport**

  1356-byte RTP packets (7 TS packets plus RTP/UDP/IPv4 headers)

  20% additional bandwidth consumption for bursting

  500 ms loss-repair buffer in each IP STB

Short-GoP Results

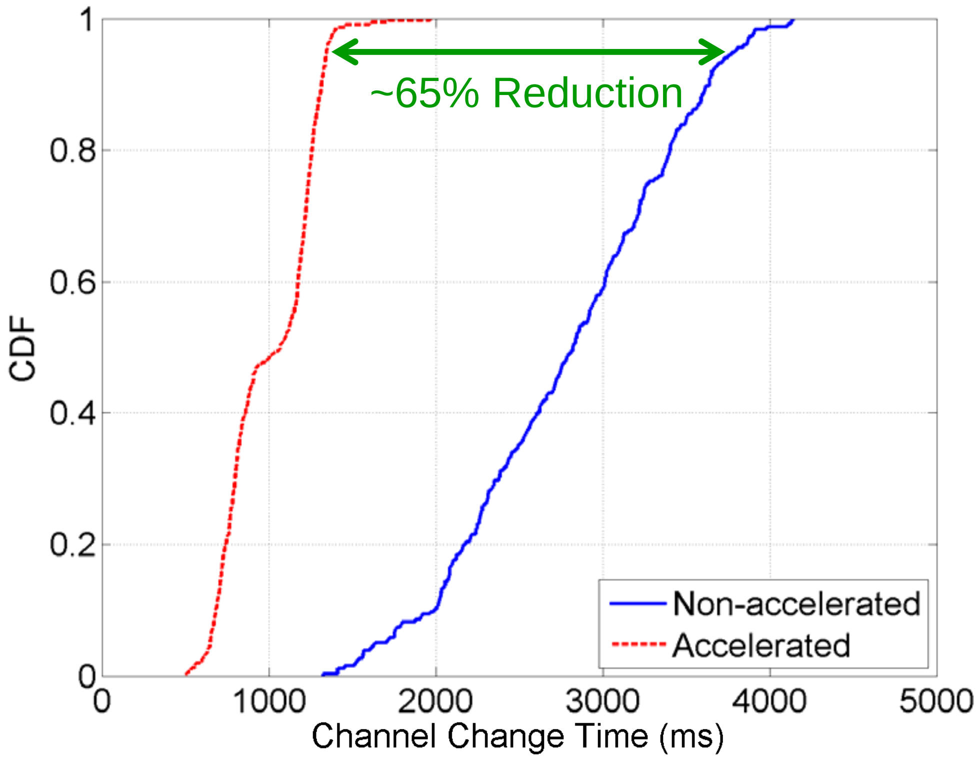

	Min	Mean	Std	95th	99th	Max
Non-accelerated	1323	2785	645	3788	4101	4140
Accelerated	501	1009	260	1345	1457	1965

Long-GoP Results

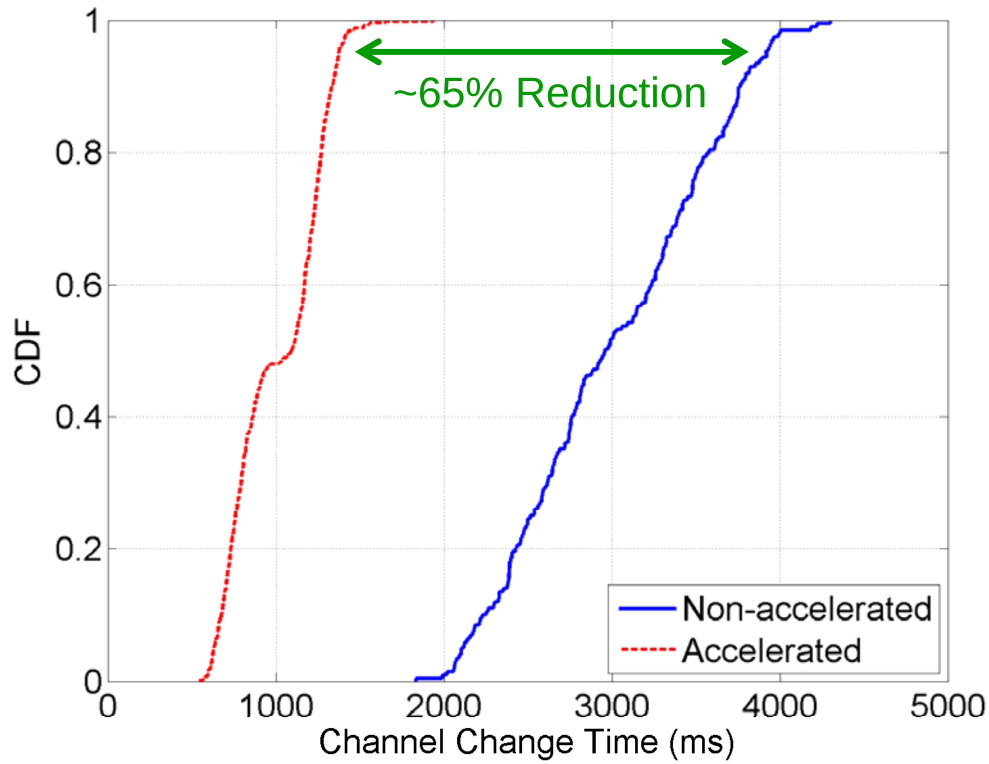

	Min	Mean	Std	95th	99th	Max
Non-accelerated	1831	3005	575	3920	4201	4300
Accelerated	536	1013	265	1377	1521	1937

When Acceleration is Disabled

Channel change times depend on

- The time of zapping request
- The frequency of the TSRAP points

→ **Inconsistent zapping times**

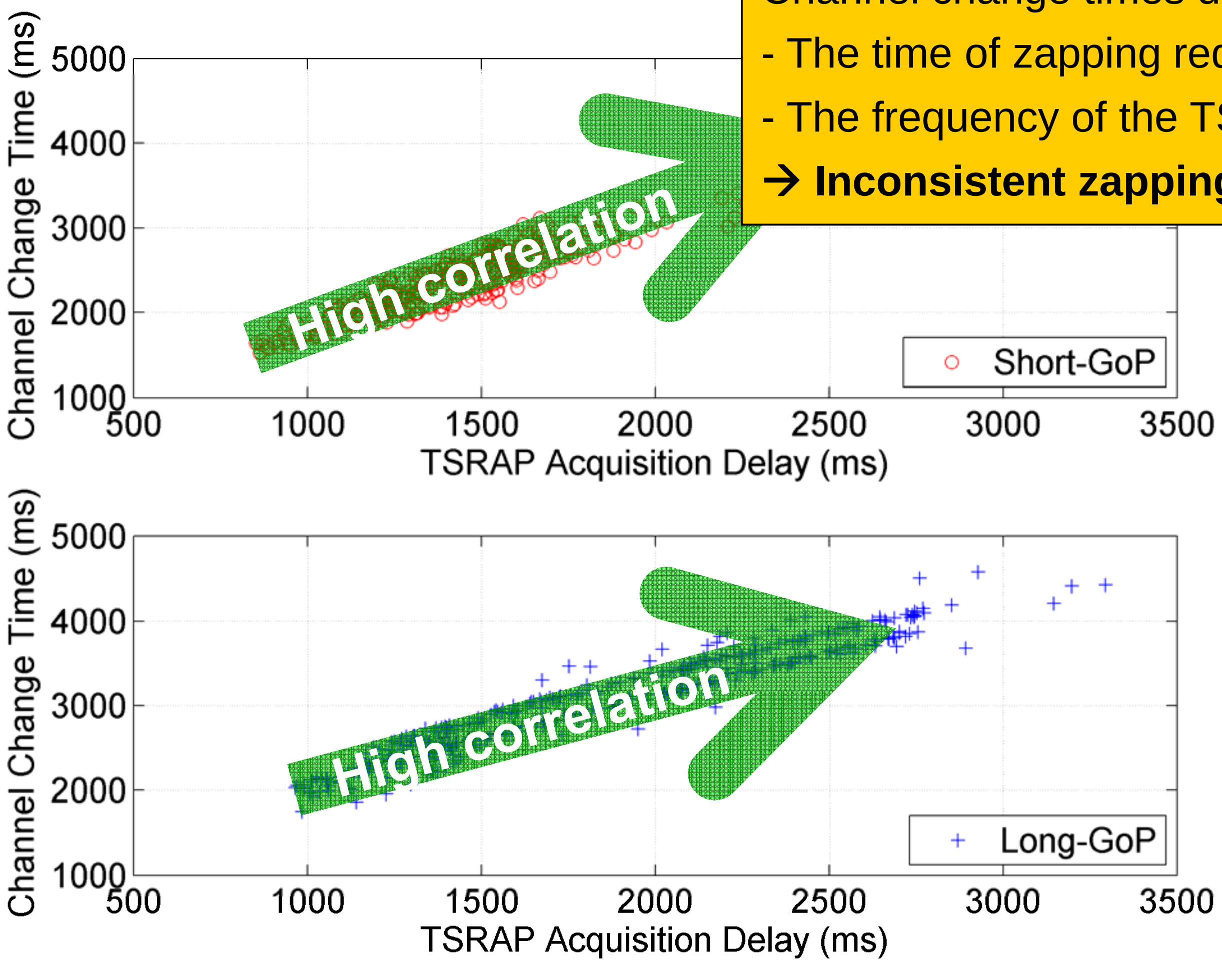

TSRAP Acquisition Delay: Time for IP STB to receive all TS-related information

When Acceleration is Enabled

Loss-repair buffer size → Loss-repair buffer size + TSRAP period

Channel change times **DO NOT** depend on

- The time of zapping request
- The frequency of the TSRAP points

→ **Consistent zapping times**

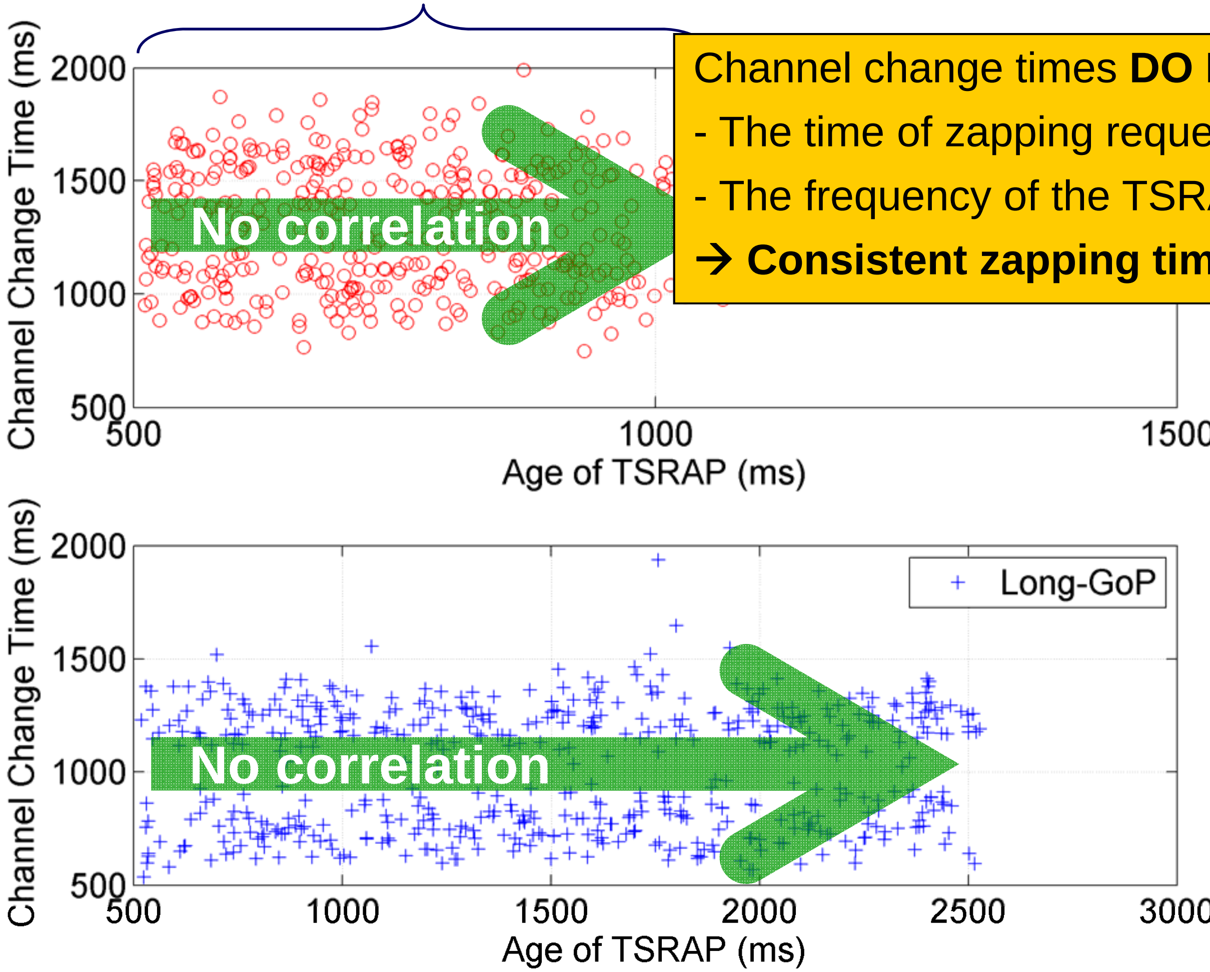

Age of TSRAP: Denotes how far TSRAP is behind multicast session when burst starts

VQE QoS/QoE Monitoring

- VQE-S collects RTCP reports
- Exporter function outputs the reports to video management application
- Management application
  - Collects raw data from exporter
  - Organizes database
  - Conducts data analysis, trends
  - Create alerts
- Management application supports standards-based north-bound interfaces
- Reports and analysis can be granular to
  - Regions
  - Edge routers
  - DSLAMs
  - Access lines
  - Home gateways
  - Settops
- Settops can support RTCP reporting and TR-069 (or TR-135) concurrently

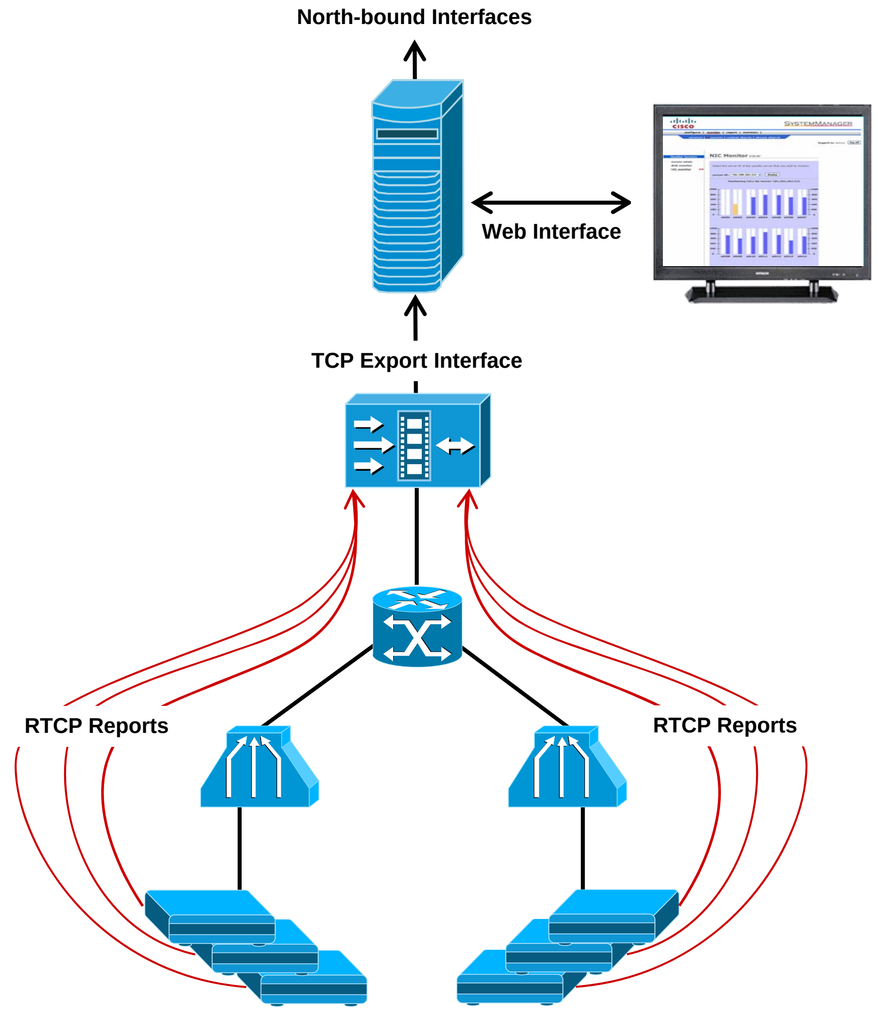

RTCP Sender/Receiver/Extended Reports

- RTCP Sender Reports provide info on data sent recently
  - Wallclock time and the corresponding RTP timestamp
  - Total number of packets/bytes sent
- RTCP Receiver Reports summarize the reception quality
  - Timestamp of (and delay from) the last received sender report
  - Highest sequence number seen so far
  - Number and fraction of the lost RTP packets
  - Estimate of the interarrival jitter
- RTCP Extended Reports (XR) can provide
  - Detailed transport-level stats and application-specific information about the RTP transport
  - Several advantages over traditional and proprietary monitoring solutions
- RTCP XR framework is easily extensible to report on
  - Packet-level loss events, loss patterns, mean time between losses, loss durations, etc.
    - → Correlation engines identify, characterize and isolate the problems
  - Audiovisual reception quality
  - Effectiveness of the loss-repair methods
    - → Loss-repair methods can be adapted and improved per network conditions
  - Effectiveness of channel change acceleration

RTCP XR Example: Loss RLE Reports

http://tools.ietf.org/html/draft-ietf-avt-post-repair-rtcp-xr

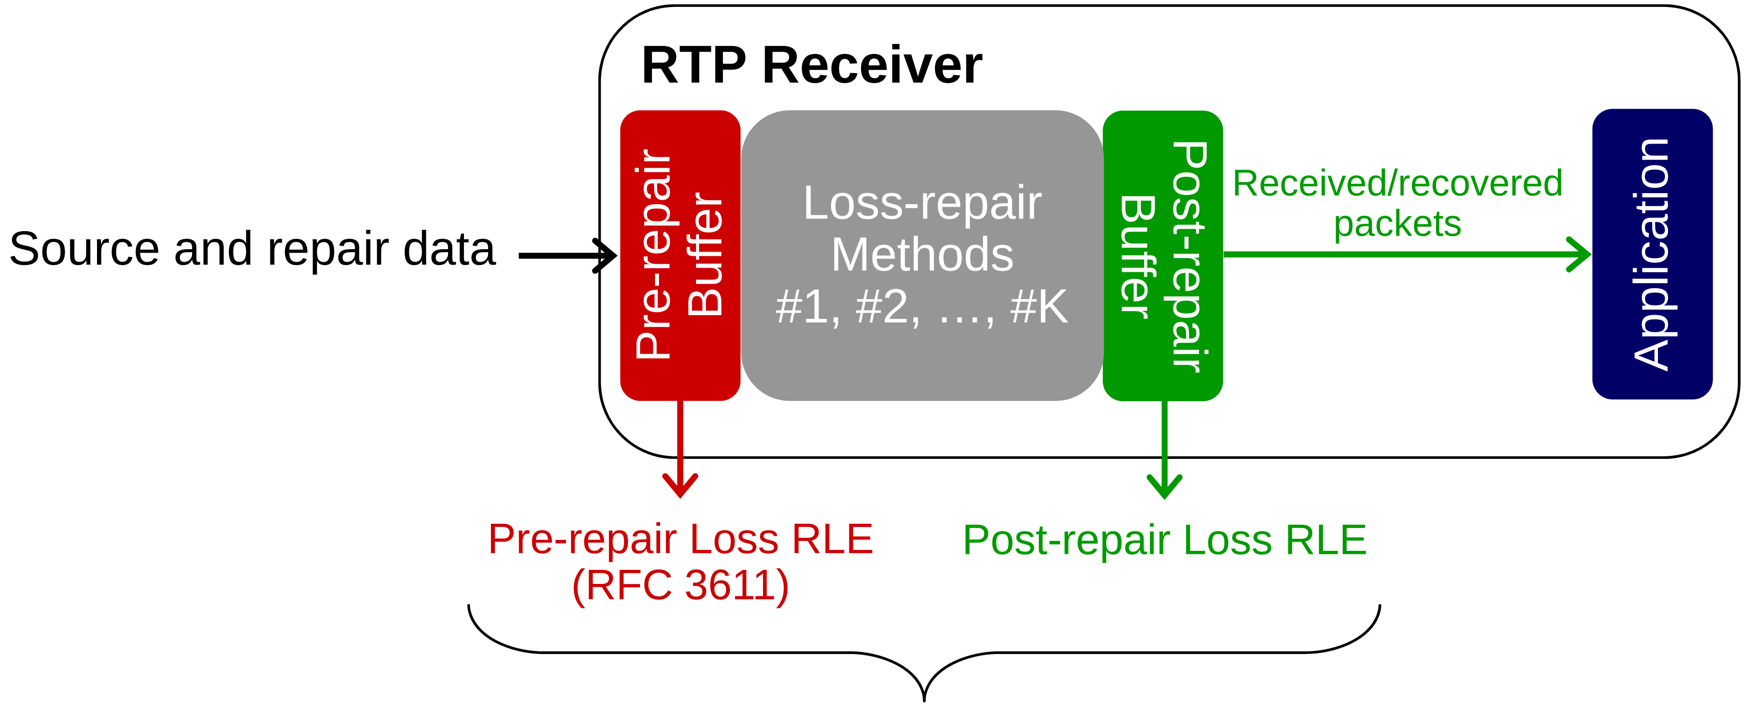

The difference tells us the aggregated performance of the loss-repair methods

Open Source Implementation for VQE Clients

- Client-side implementation is available as open source:

Documentation

http://www.cisco.com/en/US/docs/video/cds/cda/vqe/3_4/user/guide/ch1_over.html

FTP Access

ftp://ftpeng.cisco.com/ftp/vqec/

VQE – Summary

- Designed with both video and network considerations
  - Scalability
  - CAC and QoS
  - Multicast
  - High availability
- Open, standards-based solution
  - Highly extensible
  - Better interoperability
- Offers hybrid loss-repair and rapid channel change solutions
  - Improves customer satisfaction
  - Expands the IPTV coverage area
- Provides end-to-end monitoring capability of individual STBs
  - Reduces costly help-desk calls and truck rolls
  - Helps isolate the source of the problem

Selected Reading

- Visit http://ali.begen.net for our most recent papers and IETF drafts
- Check out the recent special issues/sessions in
  - IEEE Communications Magazine (Multiple issues in 2008)
  - IEEE Internet Computing (May 2009)
  - IEEE Trans. Broadcasting (June 2009)
  - IEEE CCNC 2008-2010
- Other Reading:
  - Light Reading: IPTV & Digital Video QoE: Test & Measurement Update
    - http://www.lightreading.com/insider/details.asp?sku_id=2382&skuitem_itemid=1181
  - Light Reading: Cisco Put to the Video Test
    - http://www.lightreading.com/document.asp?doc_id=177692&site=cdn
  - EANTC Experience Provider Mega Test
    - http://www.cisco.com/en/US/solutions/ns341/eantc_megatest_results.html

Questions

The answer to life's problems aren't at the bottom of a bottle, they're on TV!

Homer Simpson

