Improving the Efficiency of Streaming over Satellite Networks

by using Unequal Error Protection of mABR streams

June 24, 2025

Yuriy Reznik
Streaming Labs, Ltd
yreznik@streaminglabs.com

Satellite-Assisted Media Delivery

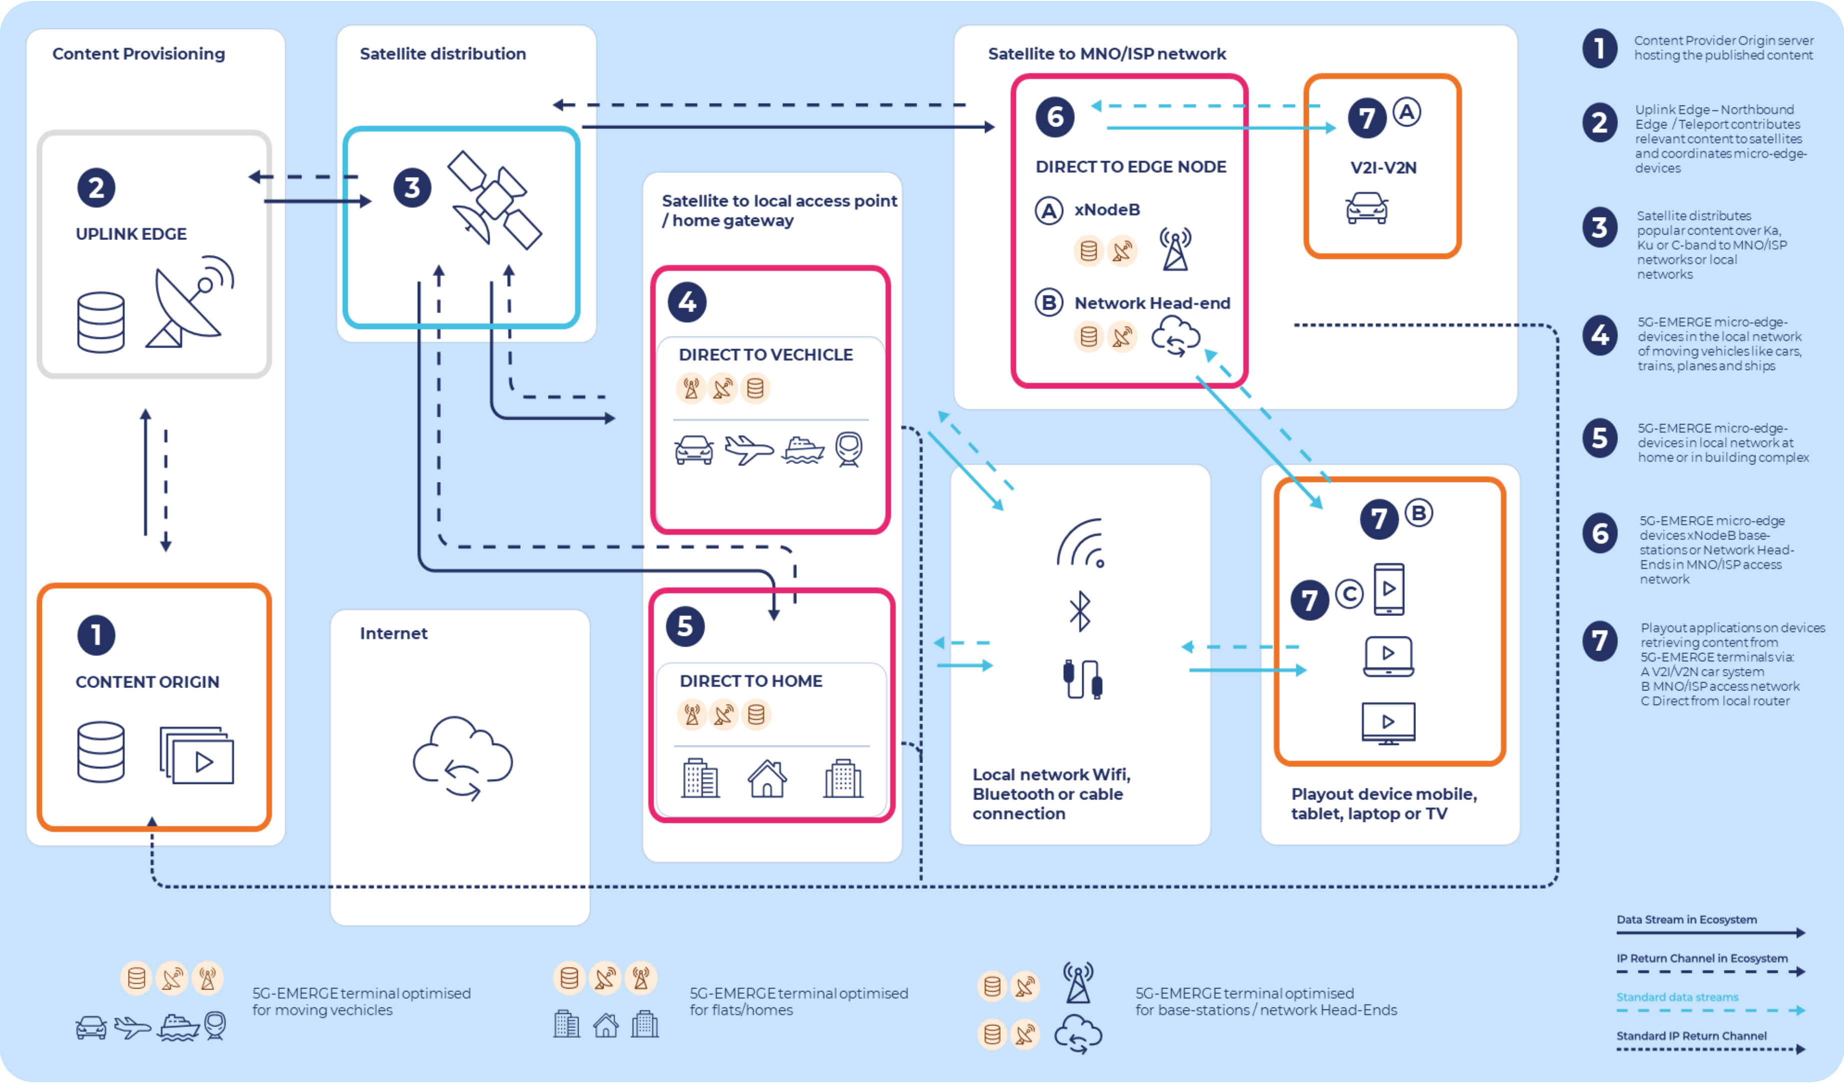

5G-EMERGE

Website
https://www.5g-emerge.com/

Members:
https://www.5g-emerge.com/consortium

White paper:
https://www.5g-emerge.com/resources

The system

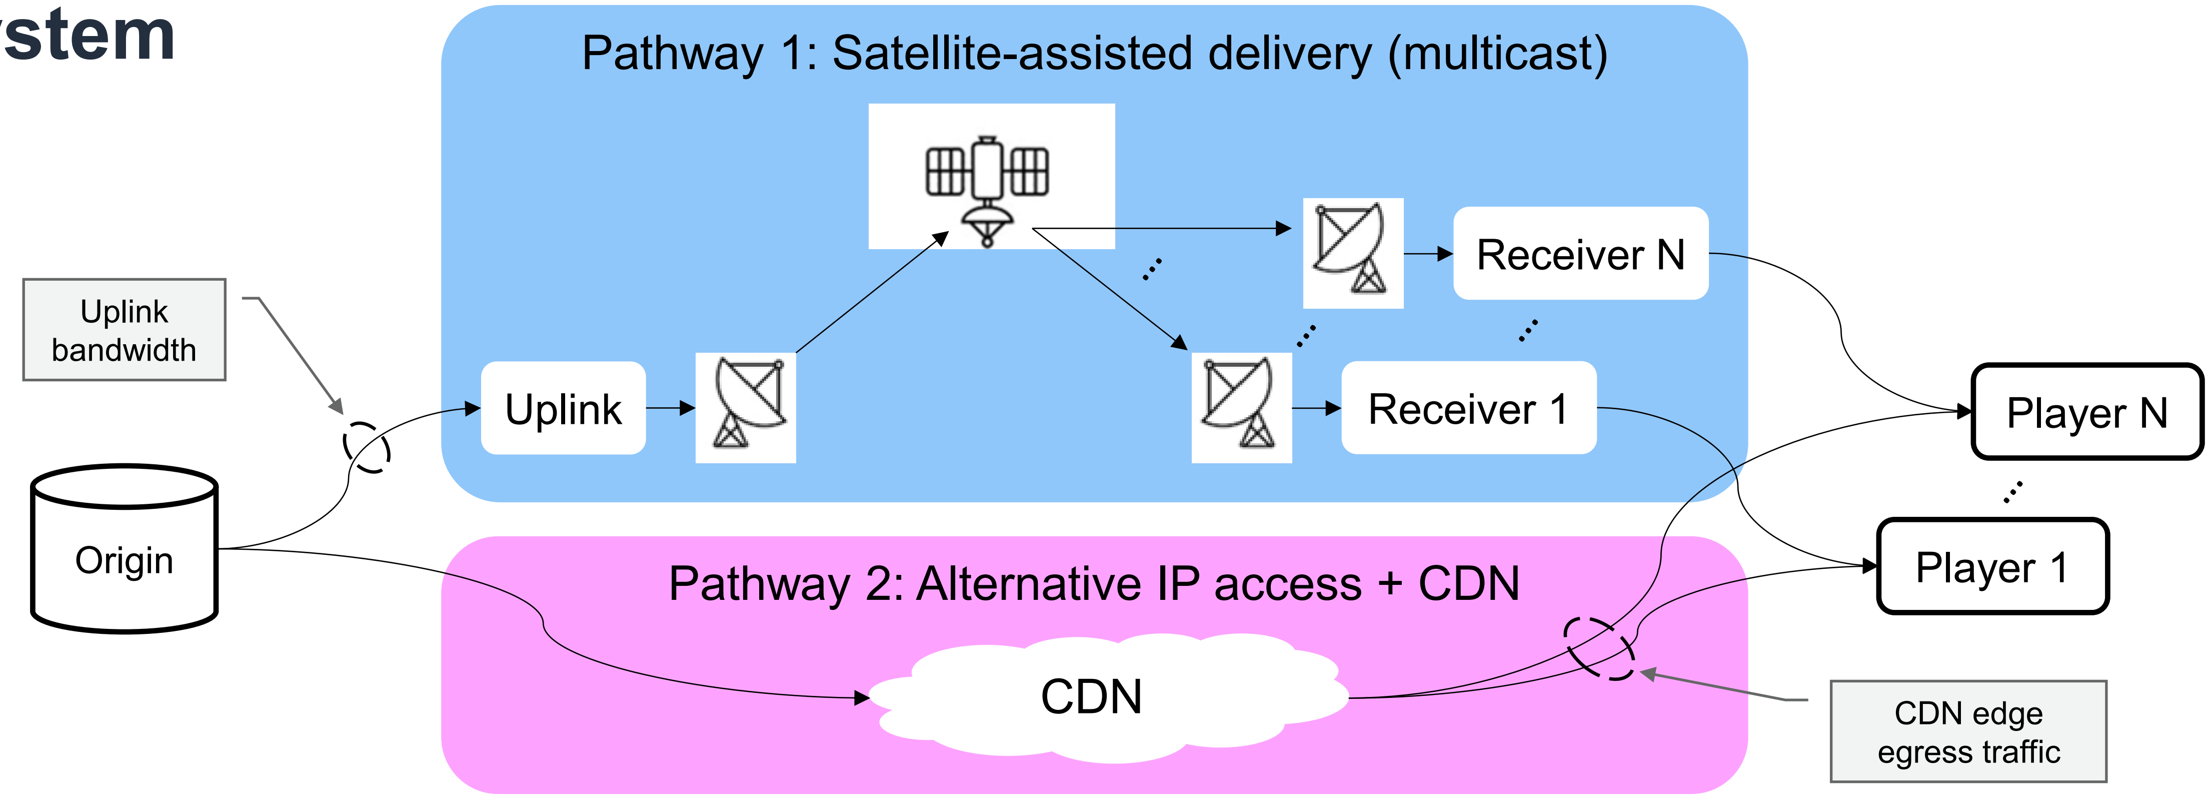

Key situations / questions:

- Pathway 2 is available ➔ How to split traffic over satellite vs CDN?
- Pathway 2 is not available ➔ How to maximize the effectiveness of satellite chain?

Cases when alternative delivery pathway is available

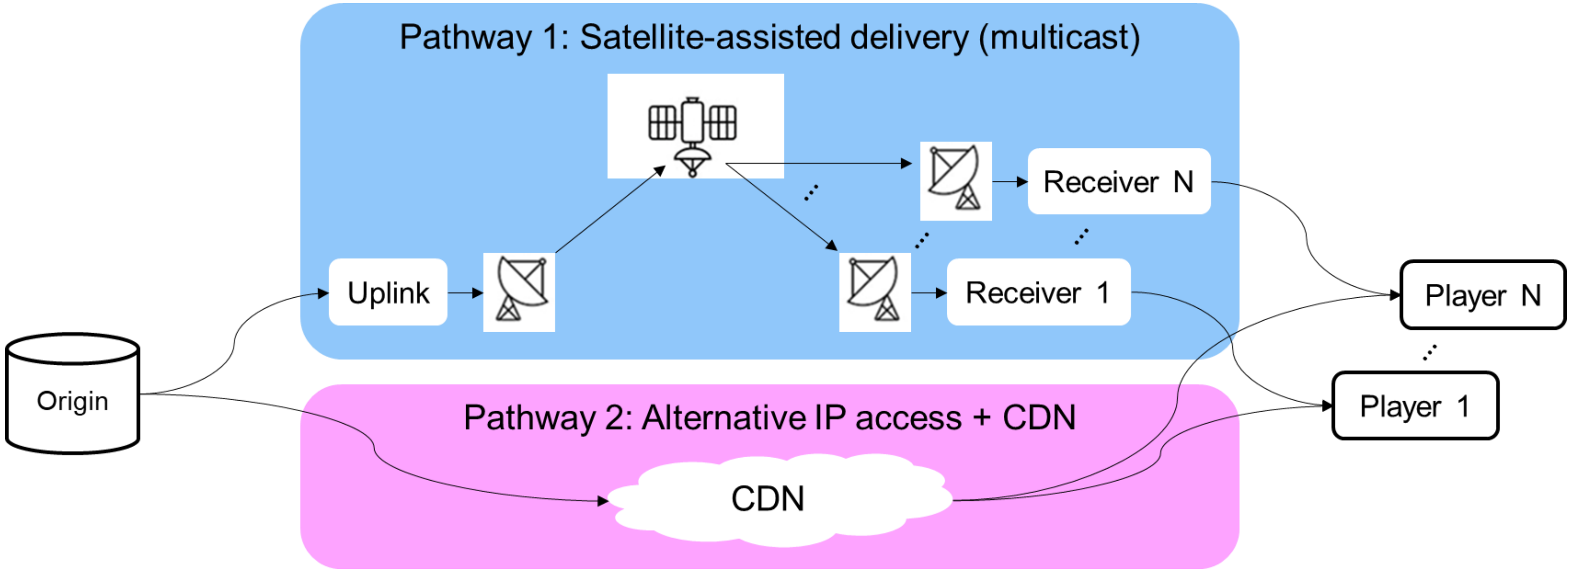

Problem 1: Splitting the traffic

System:

- Live channels 1..X
- Each channel is encoded as 1 stream
- All streams have the same bitrate R

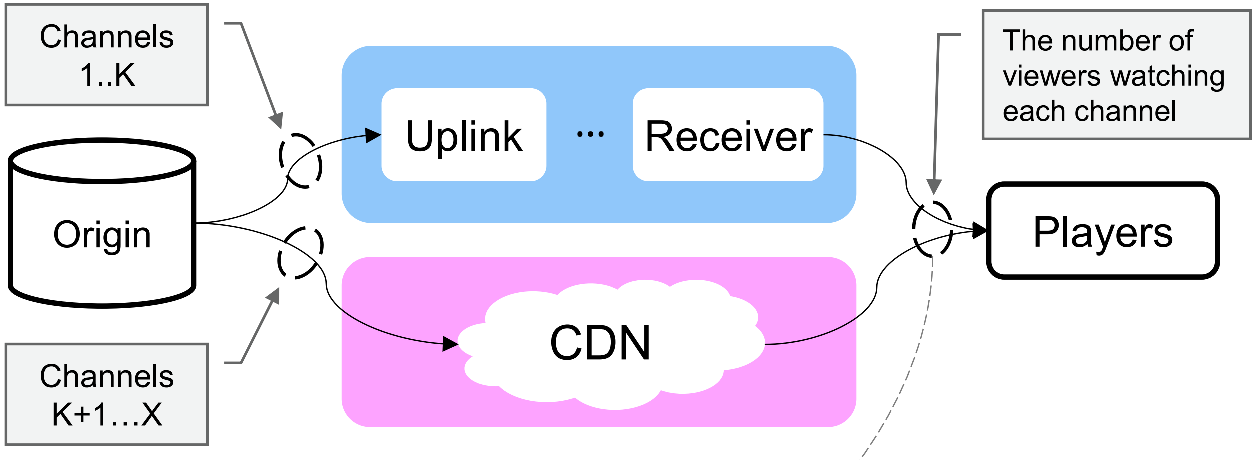

Content popularity model:

- $p(x) = \frac{x^{-\alpha}}{\zeta(\alpha)}$ – Zeta distribution, α – parameter ($\alpha > 1$)
- $p(x \leq K) = \sum_{x=1}^{K} p(x) \sim 1 - \frac{K^{1-\alpha}}{(\alpha-1)\zeta(\alpha)}\left(1 + O\left(\frac{1}{K}\right)\right)$
  - – fraction of traffic absorbed by top K channels
  - – these channels will be delivered over the satellite

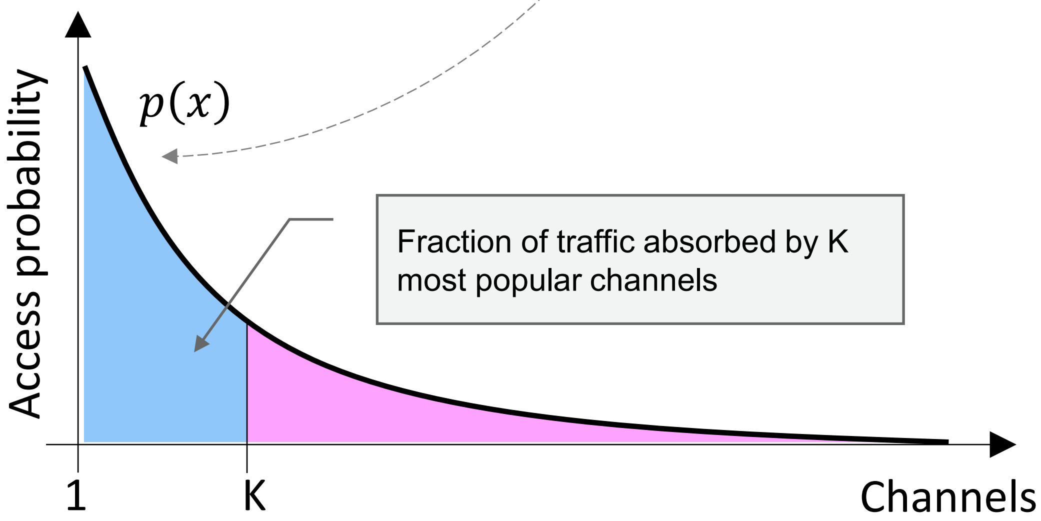

Problem 1: Splitting the traffic

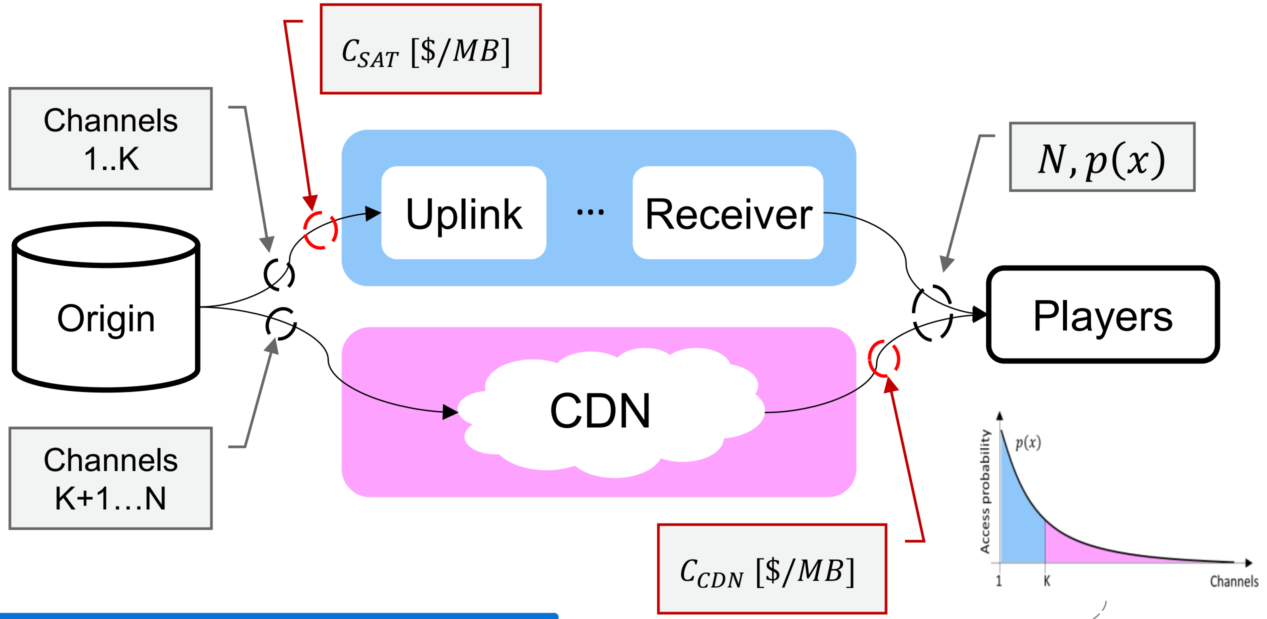

System parameters

- K – channels allocated to satellite
- N – the total number of concurrent viewers
- T – observation period [sec]
- R – rate of each channel [Mbps]
- C_{SAT} – cost of satellite bandwidth [$/MB]
- C_{CDN} – cost of CDN bandwidth [$/MB]

Total delivery cost

- $C_{total} = R \cdot T \cdot \left(C_{SAT} \cdot K + C_{CDN} \cdot N \cdot p(x > K)\right) \approx \boxed{R \cdot T \cdot \left(C_{SAT} \cdot K + C_{CDN} \cdot N \cdot \frac{K^{1-\alpha}}{(\alpha-1)\zeta(\alpha)}\right)}$

Tail of ζ - distribution

Minimum w.r.t. K

- $\frac{d}{dK} C_{total}(K) = 0 \quad \Rightarrow \quad \boxed{K_{opt} \approx \left(N \cdot \frac{C_{CDN}}{C_{SAT}} \cdot \frac{1}{\zeta(\alpha)}\right)^{\frac{1}{\alpha}}} \quad \Rightarrow \quad \left[\begin{array}{lll} N \uparrow & \Rightarrow & K_{opt} \uparrow \\ \frac{C_{CDN}}{C_{SAT}} \uparrow & \Rightarrow & K_{opt} \uparrow \\ \alpha \uparrow & \Rightarrow & K_{opt} \downarrow \end{array}\right.$

Break-even point

- $K_{opt} \geq 1 \quad \Rightarrow \quad \boxed{C_{SAT} \leq N \cdot C_{CDN} \cdot \frac{1}{\zeta(\alpha)}}$

Problem 1a: Splitting of traffic with 2 renditions (e.g. UHD & HD)

Extra parameters

- Rendition bitrates: R_1, R_2
- Rendition relative pull-rates: $q_1, q_2 \ \ (q_1 + q_2 = 1)$

Probabilities and expected rates of each stream

- x – channel id, j – rendition id
- $p(x, j) = p(x) \cdot q_j$ – probability of an (x,j) stream
- $R(x, j) = p(x, j) \cdot R_j$ – effective bitrate of (x,j) stream in the overall composition of all traffic delivered

New sorting order

- To identify which channels/renditions should go over satellite we now need to sort all steams according to their effective bitrates

$$R(x_1, j_1) \leq R(x_2, j_2) \leq \ \dots \leq R(x_X, j_X)$$

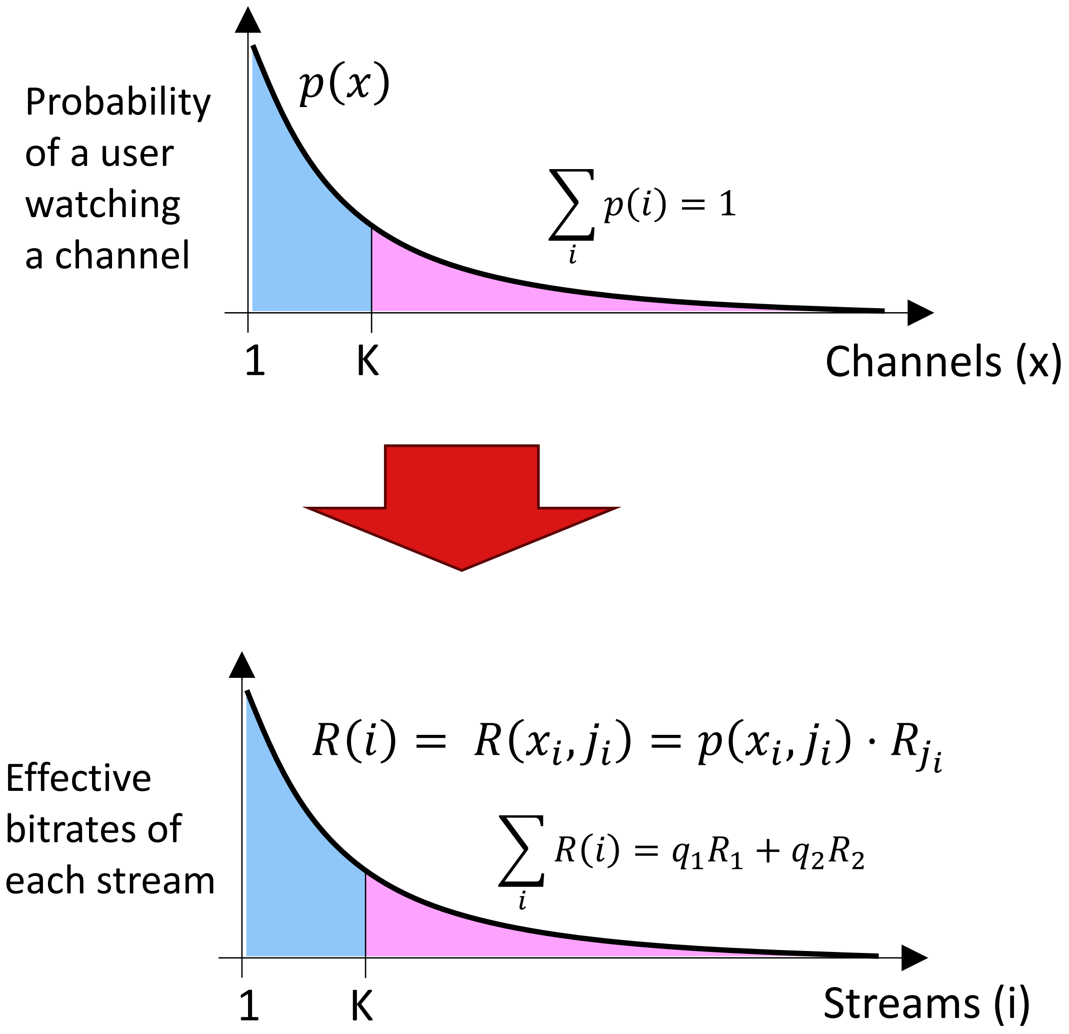

Problem 1a: Splitting of traffic with 2 renditions

Sorting streams by effective bitrates:

- R_1, R_2 – bitrates; q_1, q_2 – load probabilities; x – channel #
- Sorted table for a case when: $q_1 \cdot R_1 > q_2 \cdot R_2$:

i	Stream params	Effective bitrate	Comment
1	$x = 1, j = 1$	$R(1{,}1) = p(1) \cdot q_1 \cdot R_1$	
…	…	…	
m	$x = m, j = 1$	$R(m, 1) = p(m) \cdot q_1 \cdot R_1$	m is largest int. s.t. $R(m, 1) <$ $R(1{,}2)$
$m + 1$	$x = 1, j = 2$	$R(1{,}2) = p(1) \cdot q_2 \cdot R_2$	
$m + 2$	$x = m + 1, j = 1$	$R(m + 1{,}1) = p(m + 1) \cdot q_1 \cdot R_1$	
…	…	…	
$m + k$	$x = m + \mathrm{k}, j = 1$	$R(m + \mathrm{k}, 1) = p(m + \mathrm{k}) \cdot q_1 \cdot R_1$	$m + k \approx 2m$ is largest int. s.t. $R(m + \mathrm{k}, 1) <$ $R(2{,}2)$
$m + k + 1$	$x = 2, j = 2$	$R(2{,}2) = p(2) \cdot q_2 \cdot R_2$	
$m + k + 2$	$x = m + \mathrm{k} + 1, j = 1$	$R(m + \mathrm{k}, 1) = p(m + \mathrm{k}) \cdot q_1 \cdot R_1$	
…	…	…	

Observations:

- Renditions #2 become injected with a step size $m \approx \left(\frac{q_1 \cdot R_1}{q_2 \cdot R_2}\right)^{\frac{1}{\alpha}}$
- $R(i \leq K) = \sum_{i \leq K} R(x_i, j_i) = \sum_{x \leq K - K/m} R(x, 1) + \sum_{x \leq K/m} R(x, 2)$

$$\approx \left(1 - \frac{(K-K/m)^{1-\alpha}}{(\alpha-1)\zeta(\alpha)}\right) q_1 \cdot R_1 + \left(1 - \frac{(K/m)^{1-\alpha}}{(\alpha-1)\zeta(\alpha)}\right) q_2 \cdot R_2.$$

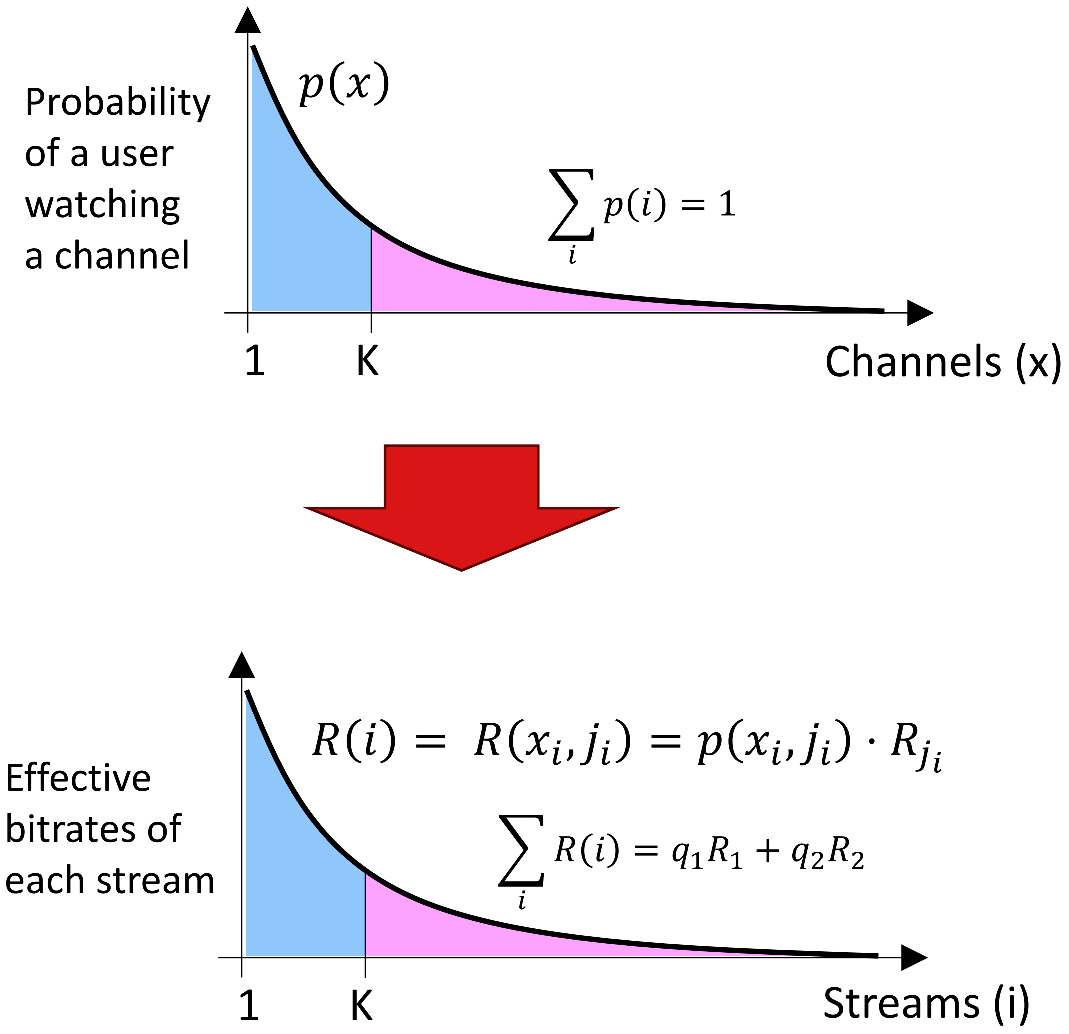

Problem 1a: Splitting of traffic with 2 renditions

System parameters

- K – channels allocated to satellite
- N – the total number of concurrent viewers
- T – observation period [sec]
- R_1, R_2 – bitrates of each renditions [Mbps]
- q_1, q_2 – pull probabilities of each renditions
- C_{SAT} – cost of satellite bandwidth [$/MB]
- C_{CDN} – cost of CDN bandwidth [$/MB]

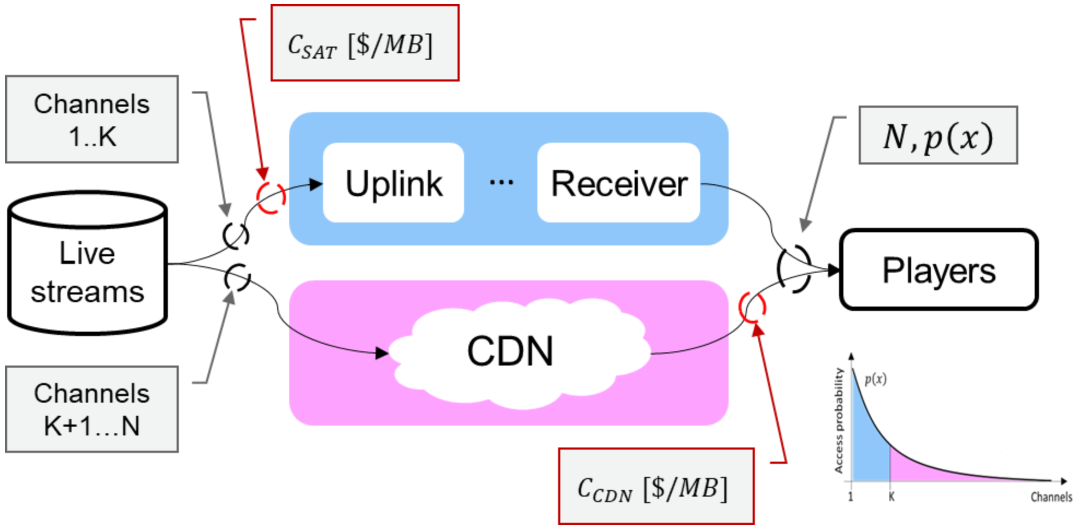

Total delivery cost

- $$C_{total} \approx T \cdot \left(C_{SAT} \left(\left(K - \frac{K}{m} \right) R_1 + \frac{K}{m} R_2 \right) + C_{CDN}\, N \left(\frac{(K-K/m)^{1-\alpha}}{(\alpha-1)\zeta(\alpha)} q_1 R_1 + \frac{(K/m)^{1-\alpha}}{(\alpha-1)\zeta(\alpha)} q_2 R_2 \right) \right), \qquad m \approx \left(\frac{q_1 \cdot R_1}{q_2 \cdot R_2} \right)^{\frac{1}{\alpha}}$$

Minimum w.r.t. K

- $$\frac{d}{dK} C_{total}(K) = 0 \quad \Rightarrow \quad \boxed{K_{opt} \approx \left(N \cdot \frac{C_{CDN}}{C_{SAT}} \cdot \frac{\left(1-\frac{1}{m}\right)^{1-\alpha} q_1 R_1 + \left(\frac{1}{m}\right)^{1-\alpha} q_2 R_2}{\left(1-\frac{1}{m}\right) R_1 + \frac{1}{m} R_2} \cdot \frac{1}{\zeta(\alpha)} \right)^{\frac{1}{\alpha}}} \quad \Rightarrow \quad \left[\begin{array}{lcl} N \uparrow & \Rightarrow & K_{opt} \uparrow \\ \frac{C_{CDN}}{C_{SAT}} \uparrow & \Rightarrow & K_{opt} \uparrow \\ \alpha \uparrow & \Rightarrow & K_{opt} \downarrow \end{array} \right.$$

Break-even point

- $$K_{opt} \geq 1 \quad \Rightarrow \quad \boxed{C_{SAT} \leq N \cdot C_{CDN} \cdot \frac{\left(1-\frac{1}{m}\right)^{1-\alpha} q_1 R_1 + \left(\frac{1}{m}\right)^{1-\alpha} q_2 R_2}{\left(1-\frac{1}{m}\right) R_1 + \frac{1}{m} R_2} \cdot \frac{1}{\zeta(\alpha)}}$$

Problem 1b: Splitting of traffic with L renditions (ABR ladder)

System parameters

- $R_1, \dots, R_L$ – bitrates of each renditions [Mbps]
- $q_1, \dots, q_L$ – pull probabilities of each renditions
- C_{SAT} – cost of satellite bandwidth [$/MB]
- C_{CDN} – cost of CDN bandwidth [$/MB]
- N – number of concurrent viewers
- T – observation time [sec]

Rendition re-ordering

- We need to re-order renditions such that

$$q_1 R_1 \geq q_2 R_2 \geq \; \dots \geq q_L R_L$$

- The rest of the processing is the same as in 2-rendtions case.

Total delivery cost

- $C_{total} \approx T \cdot \left(C_{SAT}\, K \sum_{i=1 \dots L} \phi_i R_i + C_{CDN}\, N \sum_{i=1,\dots L} K^{1-\alpha} \psi_i q_i R_i \right)$, $\quad \phi_\mathrm{i}$ and ψ_i are some funcitons of q_i, R_i, and α

Minimum w.r.t. K

- $\frac{d}{dK} C_{total}(K) = 0 \quad \Rightarrow \quad \boxed{K_{opt} \approx \left(N \cdot \frac{C_{CDN}}{C_{SAT}} \cdot \frac{\sum_{i=1,\dots L} \psi_i q_i R_i}{\sum_{i=1 \dots L} \phi_i R_i} \cdot \frac{1}{\zeta(\alpha)} \right)^{\frac{1}{\alpha}} \sim N^{\frac{1}{\alpha}}}$

Break-even point

- $K_{opt} \geq 1 \quad \Rightarrow \quad \boxed{C_{SAT} \leq N \cdot C_{CDN} \cdot \frac{\sum_{i=1,\dots L} \psi_i q_i R_i}{\sum_{i=1 \dots L} \phi_i R_i} \cdot \frac{1}{\zeta(\alpha)}}$

The effects of multiple renditions

Consider a quantity

$$\xi = \frac{\sum_{i>K} R(i)}{\sum_i R(i)} \cdot \frac{1}{\sum_{i>K} p(i)} \sim \frac{\left(\left(1-\frac{1}{m}\right)^{1-\alpha} q_1 R_1 + \left(\frac{1}{m}\right)^{1-\alpha} q_2 R_2\right)}{q_1 R_1 + q_2 R_2}$$

It defines the increase in CDN traffic due to the use of 2 renditions.

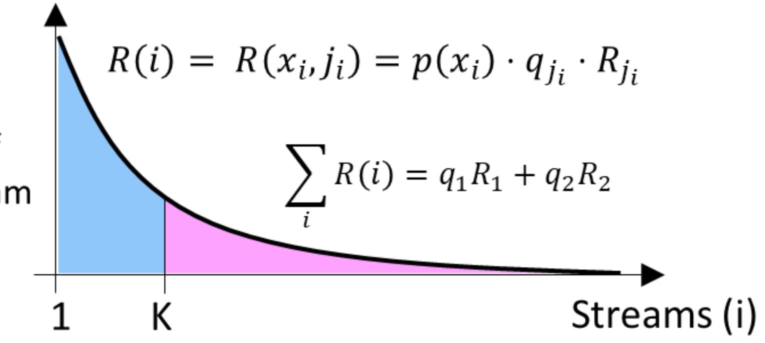

If we further introduce parameters (normalized rates):

$$\rho_1 = \frac{q_1 \cdot R_1}{q_1 \cdot R_1 + q_2 \cdot R_2}; \ \rho_2 = \frac{q_2 \cdot R_2}{q_1 \cdot R_1 + q_2 \cdot R_2} = 1 - \rho_1;$$

Then

$$\xi\,(\alpha, \rho_1) \sim \left(1-\frac{1}{m}\right)^{1-\alpha} \rho_1 + \left(\frac{1}{m}\right)^{1-\alpha} \rho_2, \quad m = \left(\frac{\rho_1}{\rho_2}\right)^{\frac{1}{\alpha}}$$

implying that ξ depends only on α and ρ_1

Observation:

$$\boxed{\rho_1 \to 1 \Rightarrow \xi \to 1\ !}$$

Plot of a function $\xi\,(\alpha, \rho_1)$

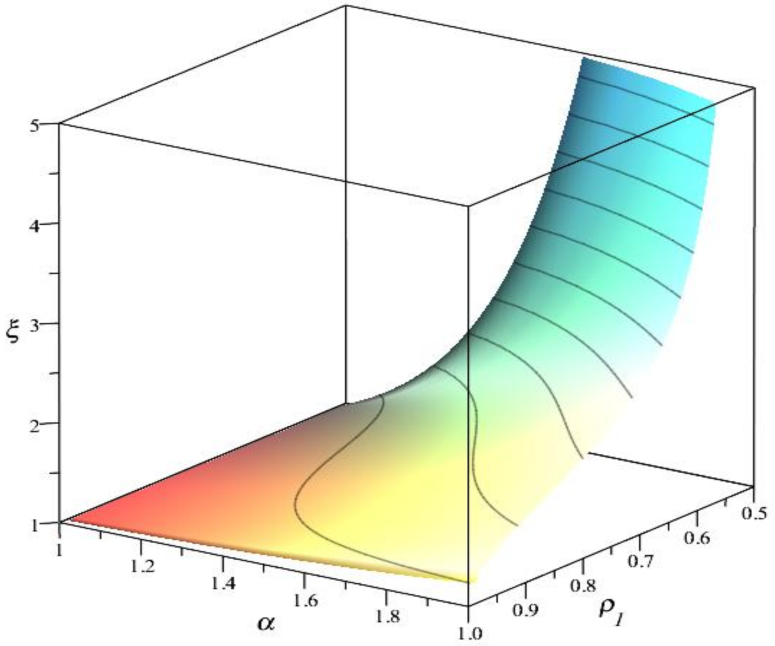

STREAMING LABS

Splitting traffic in the hybrid delivery system: the summary

General observations

- The problem is very simple
- Solvable in cases of single streams / channel
- Also solvable for cases of ABR ladders
- Provides formulae for optimal allocation of satellite capacity

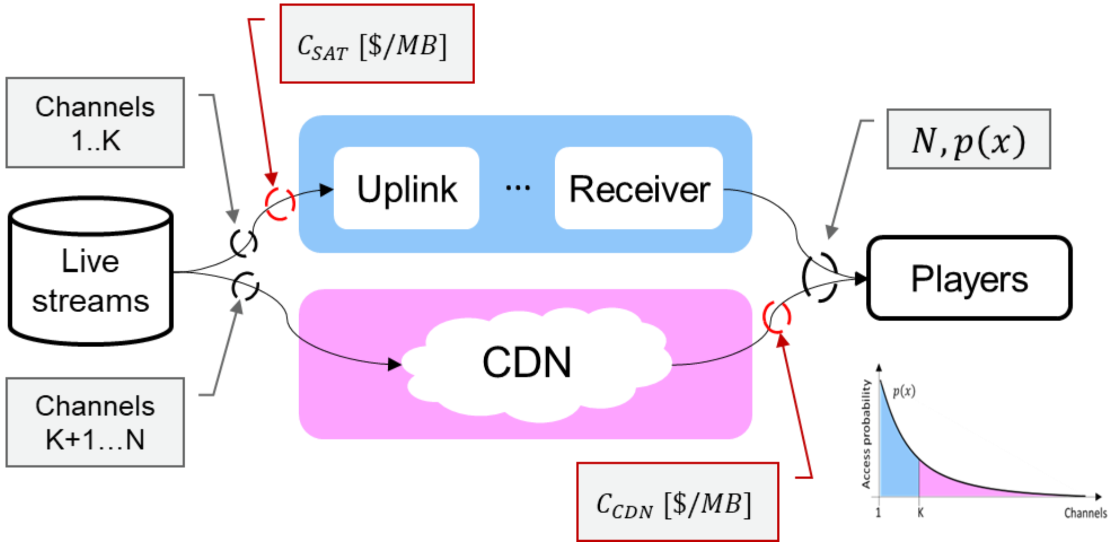

Parameters that influence the solution

- α – parameter or channel popularity distribution; the peakier it is $(\alpha \gg 1)$, the fewer channels can go over satellite
- N – number of concurrent viewers; the higher it is, the more channels will make sense to send over the satellite.
- C_{SAT}/C_{CDN} – satellite vs CDN cost ratio. The larger it is, the fewer channels it would make sense to send over satellite.
- $\rho_i = R_i q_i / \sum_i R_i q_i$ – relative effective bitrates of renditions in the ladder. The more asymmetric this distribution is, the better!

Cases when alternative delivery pathway is not available

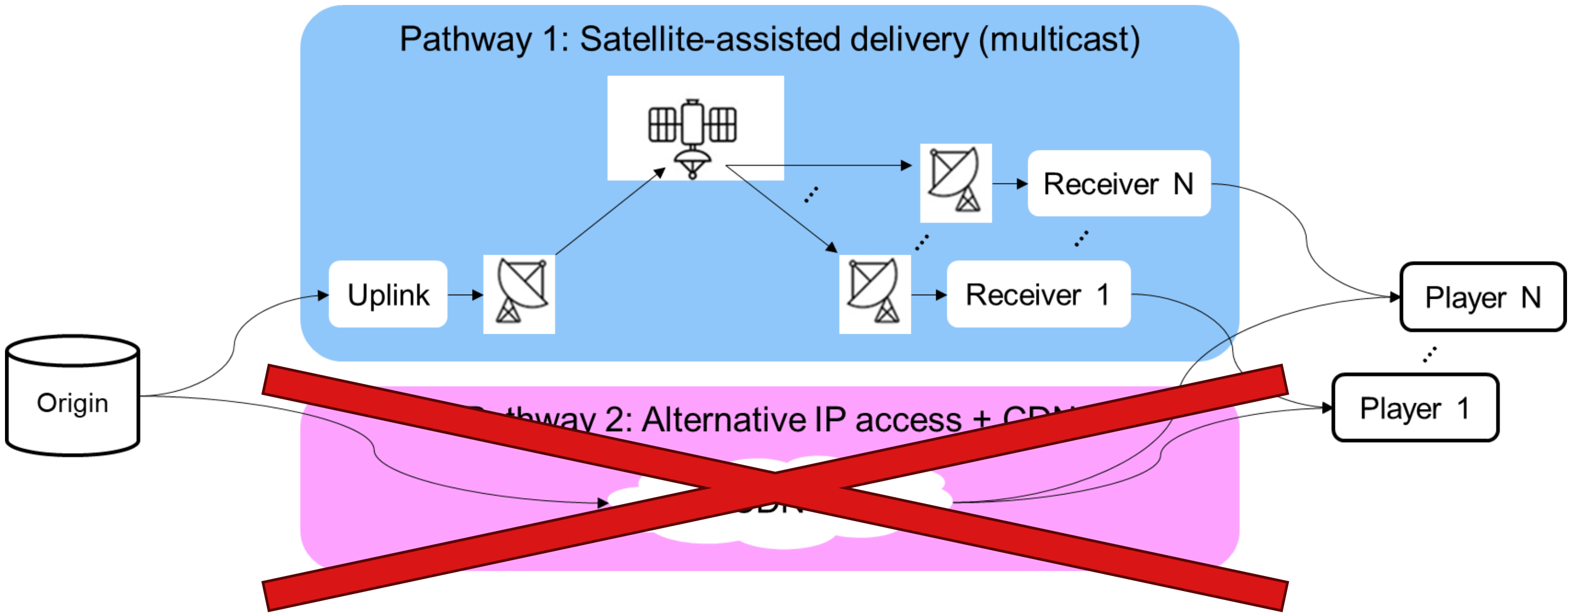

Multicast streaming over a satellite

System model

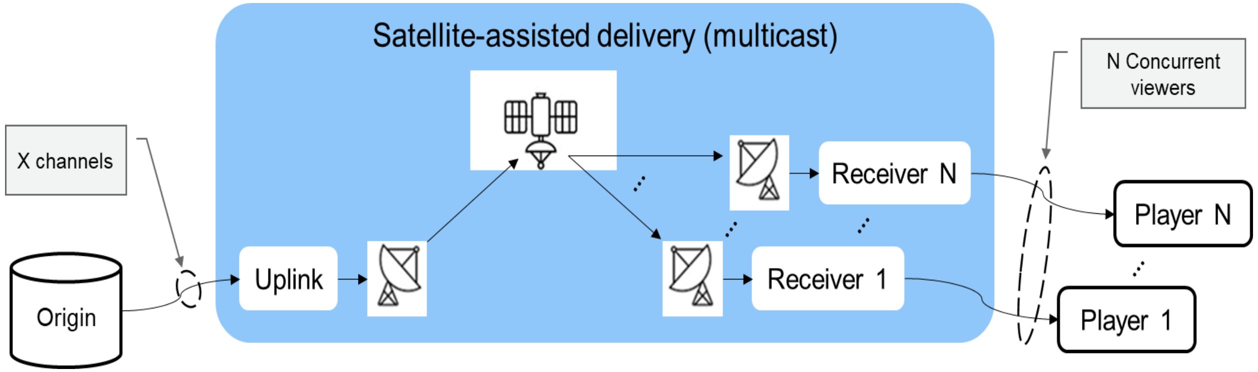

Extra details

- The delivery of ABR streaming data is typically managed by using the mABR standard. It pushes all streaming data to uplink.
- At the receiver's side, all data are placed on a local HTTP server, which HLS/DASH player then accesses for playback.
- If by-directional transmission is available, mABR standard defines mechanisms for pulling missing data and trying to repair data that were partially delivered. But it is not always available.

ETSI TS 103769 (2024) – MABR Standard

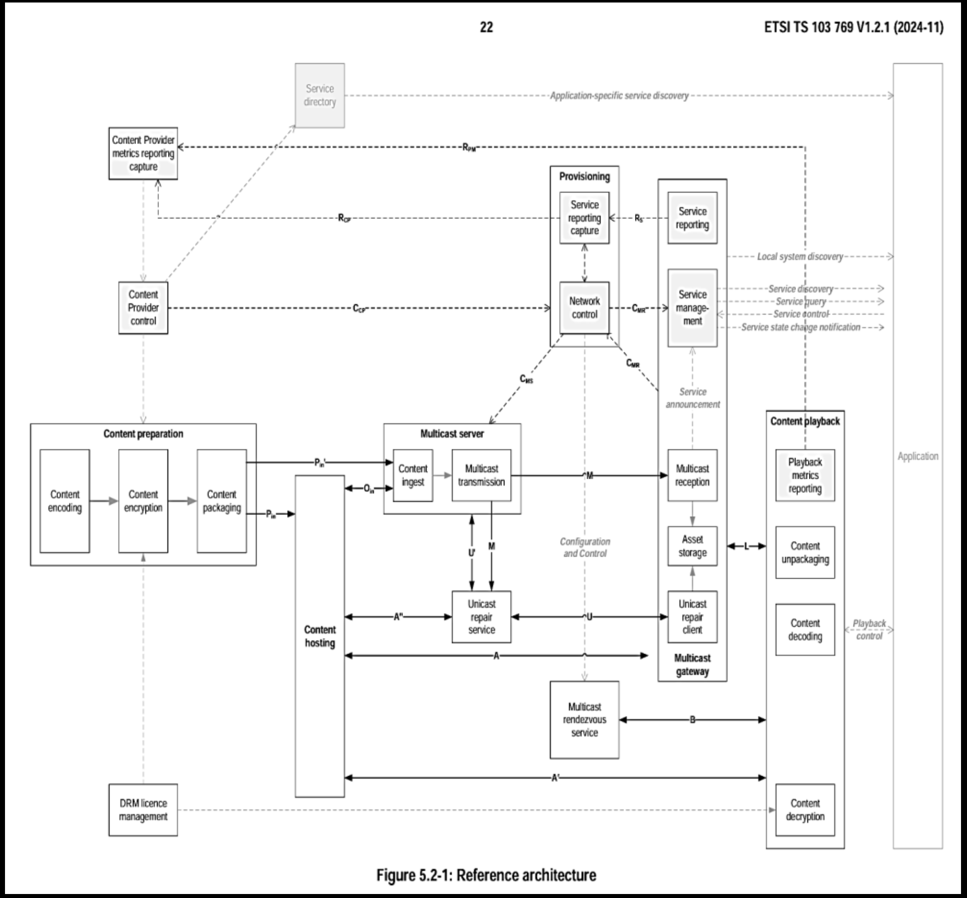

Unidirectional data delivery. FECs, AL-FECs.

General notes

- Normally, satellites add physical layer FECs and modulate transmitter power as primary mechanisms ensuring data delivery. But sometimes (e.g. under heavy rains) some data may be lost.
- As an extra mechanism for <u>increasing probability</u> of data being delivered on can use AL-FECs, as supported by mABR / FLUTE.

Example

- RS code with k data packets, and n-k redundancy packets
- If at least k packets out of n are received, the data are recoverable
- This means that if we have a channel with PER = p, then the probability that a (k,n) RS code will be successful in recovering data packets is

$$P_{\text{delivered}}(k,n,p) = \Pr(t \le n-k) = \sum_{t=0}^{n-k} \binom{n}{t} p^t (1-p)^{n-t}$$

- Indeed, if $n \to \infty$, then:

$$P_{\text{delivered}}(k,n,p) \to \begin{bmatrix} 1 & if\ n-k \ge n \cdot p \\ 0 & otherwise \end{bmatrix}$$

Probability of success of RS code (n=128)

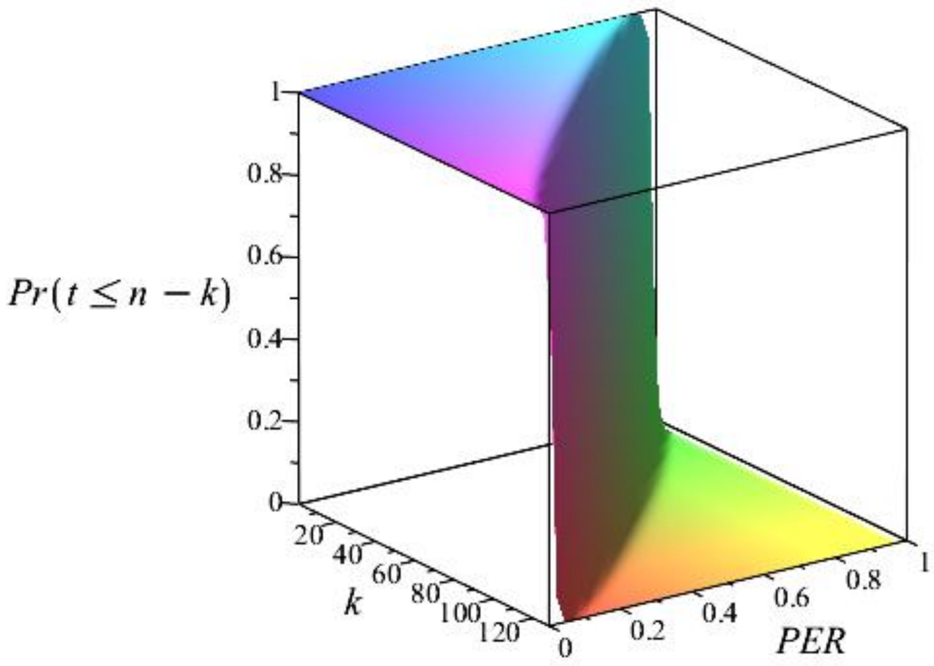

Model: memoryless channel with fixed PER

Using UEP for satellite delivery of mABR streams

The problem

- The reception of satellite signals can different across receivers
- It is heavily influenced by atmospheric conditions (rain, fog, snow) at the location of each recipient.
- Increasing power or increasing redundancy across all layers / renditions seem like brute force solutions.

The idea

- Use unequal error protection
- Assign more redundant AL-FECs at lower bitrate streams such that they can survive even the worst conditions.
- Use extra logic at edge servers to guide HLS/DASH clients toward choosing streams that are successfully downloaded.

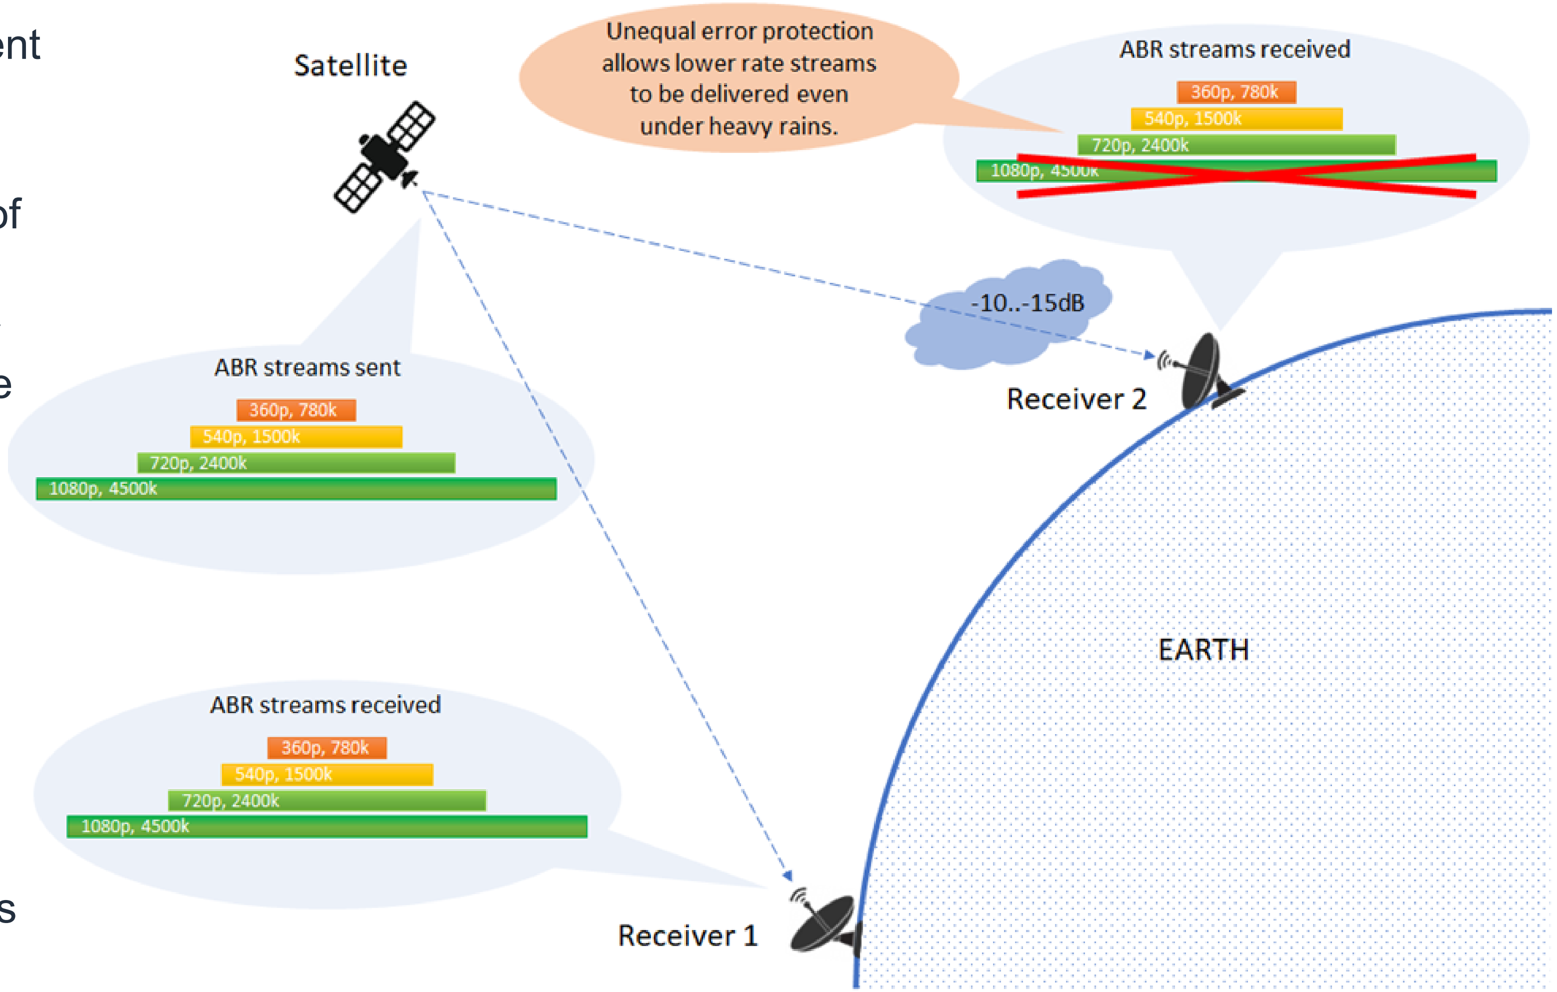

Trivialized example of using UEP

Encoding ladder

- Let's assume that we have 2 streams
- R_1, R_2 $(R_1 > R_2)$ – rendition bitrates
- Q_1, Q_2 $(Q_1 > Q_2)$ – quality estimates

AL-FECs

- $\varepsilon_1, \varepsilon_2$ – added redundancies
- $R_1(1+\varepsilon_1)$ – total rate of 1st stream
- $R_2(1+\varepsilon_2)$ – total rate of 2st stream

PER rates at receivers

- N receivers total
- K receivers have PER = p_H (high)
- $N-K$ receivers have PER = p_L (low)

Possible AL-FEC rate assignments $\left(\boldsymbol{\varepsilon_H \sim \frac{p_H}{1-p_H},\ \varepsilon_L \sim \frac{p_L}{1-p_L}}\right)$:

FEC allocation model	**High-High**	**High-Low**	**Low-High**	**Low-Low**
Redundancies	$\varepsilon_1 = \varepsilon_H;\ \ \varepsilon_2 = \varepsilon_H$	$\varepsilon_1 = \varepsilon_H;\ \ \varepsilon_2 = \varepsilon_L$	$\varepsilon_1 = \varepsilon_L;\ \ \varepsilon_2 = \varepsilon_H$	$\varepsilon_1 = \varepsilon_L;\ \ \varepsilon_2 = \varepsilon_L$
Total rate to satellite	$R_1(1+\varepsilon_H) + R_2(1+\varepsilon_H)$	$R_1(1+\varepsilon_H) + R_2(1+\varepsilon_L)$	$R_1(1+\varepsilon_L) + R_2(1+\varepsilon_H)$	$R_1(1+\varepsilon_L) + R_2(1+\varepsilon_L)$
Delivered renditions	N 1st and N 2nd	N 1st and N-K 2nd	N-K 1st and N 2nd	N-K 1st and N-K 2nd
Average quality delivered	Q_1	Q_1	$\frac{N-K}{N}Q_1 + \frac{K}{N}Q_2$	K receivers cannot play!
Tradeoff achieved	Highest bit expenditure High quality of playback	The same quality, fewer bits.	Somewhat lower quality, even fewer bits.	Not workable.

NB: Unequal bit allocation models (High-low and Low-High) have clear benefits!

Posing a more complete problem

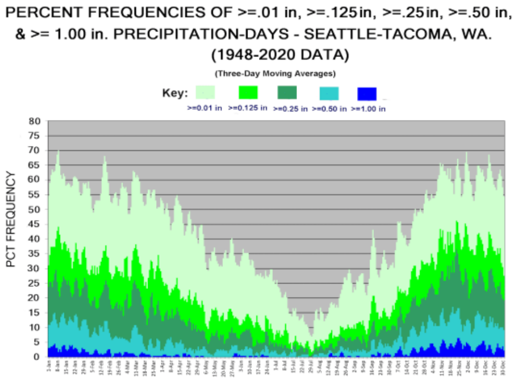

A lot of useful data, models, and statistics are available

- Per-geo location & time weather data
- Models of impact of rainfalls on received signal power level
- Models of the relationship between signal power level and packet loss

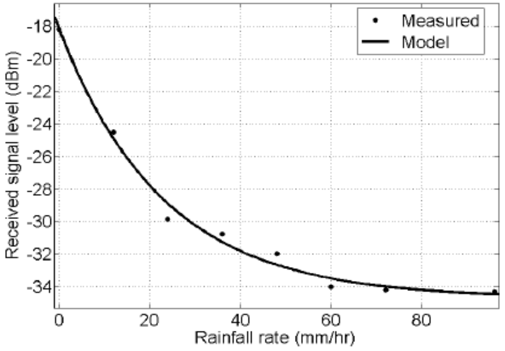

Figure 3. 4: Signal level vs. rainfall rate (measured data and model)

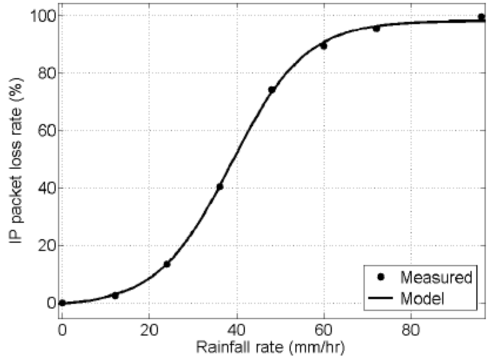

Figure 3. 6: IP packet loss-rainfall rate: measurements versus model.

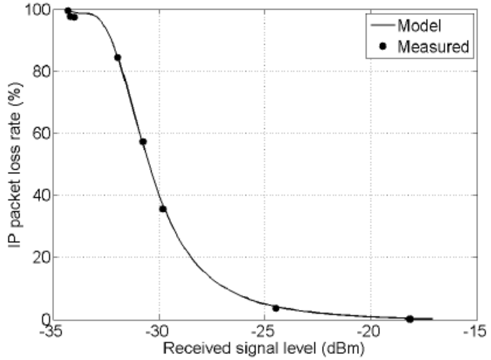

Figure 3. 7: The relationship between the received power level and IP packet loss rate.

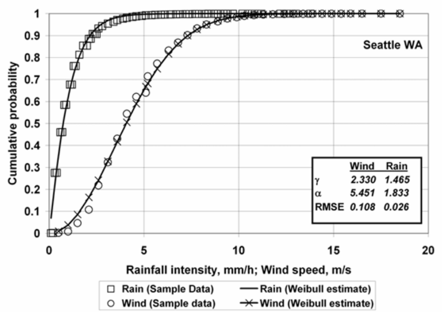

FIGURE 5 – Cumulative distribution functions for rainfall intensity (mm/h) and wind speed (m/s) for Seattle, WA, and corresponding Weibull estimates. The shape and scale parameters as well as the root mean squared error between the actual and assumed distributions are given. The period of record was 26 years.

Therefore, one can

- Obtain a distribution of packet-loss rates p across users of a specific satellite
- And pose a problem of finding an optimal rate allocation assignment $\varepsilon = [\varepsilon_1, \dots, \varepsilon_L]$ for an ABR ladder such that

$$Q_{average}(R, \varepsilon^*, p) = \max_{\substack{\varepsilon = [\varepsilon_1, \dots, \varepsilon_L] \\ R_{total} = \sum_i R_i (1+\varepsilon_i) \le R_{\max}}} Q_{average}(R, \varepsilon, p)$$

Future work

- There is clearly more work that needs to be done to produce simulations of systems using UEP for mABR delivery
- However, the sheer diversity of weather date across regions and dynamically changing nature of them suggests that this approach should be fruitful
- Effectively, with mABR and simulcasting multiple HLS / DASH streams we already build significant redundancy in the distribution. UEP + HLS/DASH selective logic offers a way of exploiting it.
- Of course, the better way of delivering videos over multicast is by using layered codecs. But that likely will need much more time to gain acceptance.

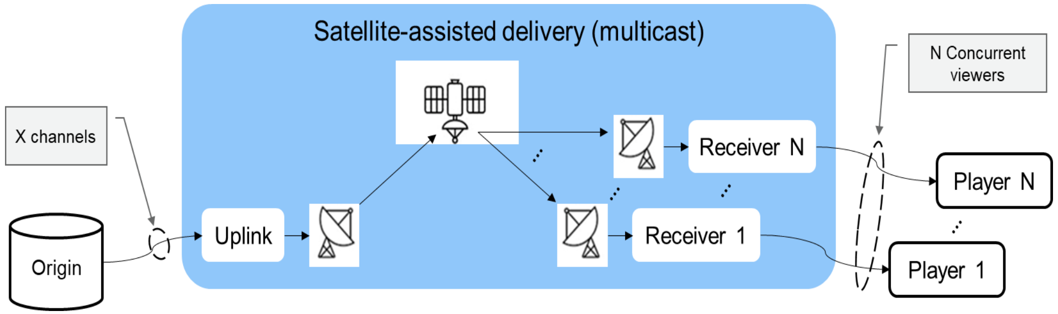

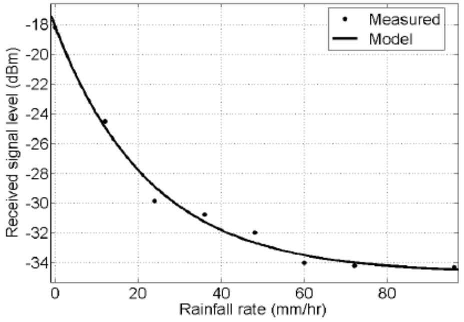

Figure 3. 4: Signal level vs. rainfall rate (measured data and model).

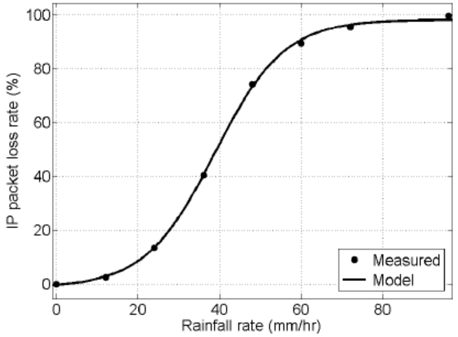

Figure 3. 6: IP packet loss-rainfall rate: measurements versus model.

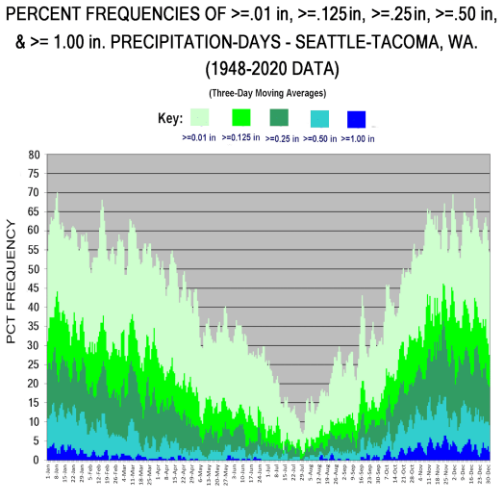

THANK YOU!

Yuriy Reznik
Streaming Labs, Ltd
yreznik@streaminglabs.com